

EXHIBIT 19

Series Editor: S. Merrill Weiss

FOCAL PRESS MEDIA TECHNOLOGY PROFESSIONAL



SECOND EDITION

Video Over IP

V, Internet Video, H.264, P2P, Web TV, and
aming: A Complete Guide to Understanding
the Technology

Wes Simpson

A **Focal Press** Book

First published 2008

This edition published 2016 by Focal Press

Published 2019 by Routledge
2 Park Square, Milton Park, Abingdon, Oxon OX14 4RN
52 Vanderbilt Avenue, New York, NY 10017

Routledge is an imprint of the Taylor & Francis Group, an informa business

© 2008 Wes Simpson.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of the publisher.

Notices

Practitioners and researchers must always rely on their own experience and knowledge in evaluating and using any information, methods, compounds, or experiments described herein. In using such information or methods they should be mindful of their own safety and the safety of others, including parties for whom they have a professional responsibility.

Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation without intent to infringe.

Library of Congress Cataloging-in-Publication Data

Simpson, Wes.

Video over IP / by Wes Simpson. – 2nd ed.

p. cm.

Includes index.

ISBN 978-0-240-81084-3 (pbk. : alk. paper) 1. Image transmission. 2. Digital video. 3. Telecommunication. I. Title.

TK5105.2.S56 2008

006.7-dc22

2008020606

British Library Cataloguing-in-Publication Data

A catalogue record for this book is available from the British Library.

ISBN 13: 978-0-240-81084-3 (pbk)

CONTENTS

Introduction to the Second Edition	xi
Introduction	xiii
Purpose of this Book	
Intended Audience	
How to Read this Book	
Video User Checklist	
Acknowledgments	
Chapter 1: Overview of Video Transport	1
Defining IPTV	
Internet Video	
Video Transport Technologies	
Telecom Networks	
The Internet	
Review and Video User Checklist	
Chapter 2: IP Video Transport Applications	29
Entertainment	
Interactive Video	
Streaming Video and Narrowcasting	
The True Meaning of "Real Time"	
Video Transport Economics	
Review and Checklist Update	
Chapter 3: Video Basics	53
Pixels, Luma, Scanning, and Chroma	
Types of Video	
Video Fields and Frames	
Types of Audio	
Other Video Services	
Video and Audio Switching	
Review and Checklist Update	

Chapter 4: Video and Audio Compression	89
Compression Overview	
MPEG Compression Technologies	
Other Compression Technologies	
Comparing Compression Technologies	
Technology Licensing Issues	
Review and Checklist Update	
Chapter 5: IP Networking Basics	137
How IP Fits In	
IP Basics	
Ethernet and IP	
Review and Checklist Update	
Chapter 6: From Video into IP Packets	155
Encapsulation	
MPEG Stream Types	
Transport Protocols	
Review and Checklist Update	
Chapter 7: IP Packet Transport	181
Transport Methods	
Transport Considerations	
Network Impairments	
Internet Transport	
Quality of Service	
Review and Checklist Update	
Chapter 8: Private Video Streaming and Media Players	219
Basic Concepts	
Streaming System Architecture	
Streaming Applications	
Technologies for Streaming	
Commercial Players	
Review and Checklist Update	
Chapter 9: Multicasting	249
Basic Concepts	
Applications	
Multicasting System Architecture	
System Impact	
Review and Checklist Update	
Chapter 10: Videoconferencing Over IP	273
Basic Concepts	
IP Videoconferencing Technologies	
Review and Checklist Update	

Chapter 11: DRM, Content Ownership, and Content Security	297
Acquiring Content	
Rights Management	
Encryption	
Encryption Systems	
Review and Checklist Update	
Chapter 12: Private Networks Secure and Transport	315
Private Networks	
Virtual Private Networks	
Review and Checklist Update	
Chapter 13: IPTV—Delivering Television to Consumers	333
Characteristics of IPTV	
Applications	
Basic Concepts	
IPTV Delivery Networks	
Technologies for IPTV	
Review and Checklist Update	
Chapter 14: Video File Transfer, Podcasting, and P2P	363
Overview	
File Transfer Applications	
File Transfer Technology	
Podcasting	
P2P File Sharing	
Review and Checklist Update	
Chapter 15: Internet Video	381
Key Attributes of Internet Video	
Internet Video Applications	
System Architecture	
Delivery Platforms and Services	
Client Software	
Review and Checklist Update	
Chapter 16: Network Administration	405
Management Tasks	
Technologies for Network Administration	
Review and Checklist Update	
Appendix A: Diffie-Hellman Key Exchange	425
Appendix B: IP Video User Checklist	429
Glossary	443
Index	465

1

OVERVIEW OF VIDEO TRANSPORT

Transporting video signals is a round-the-clock business throughout the world today. Whether for entertainment, education, or personal communication, we now live in a world where we are exposed to video content in many forms. The scale of the technologies and systems that are used to gather and deliver all of this content are amazing. For example, the 2006 FIFA World Cup™ tournament had a cumulative television audience of over 26 billion viewers during the 30-day tournament. But Internet video is rapidly catching up—a January 17, 2008 press release from comScore reported that Americans watched 9.5 billion Internet videos in November 2007, with the average viewer watching 69 videos at an average length of 2.8 minutes, representing an increase of 29 percent since the beginning of 2007.

As Internet Protocol (IP) technologies for video transport continue to mature, more of the video delivery process will take place over IP networks. This change will ultimately include all phases of video content creation and delivery, beginning at the video camera and ending at a home viewer's video display. As we will see throughout

this book, many different areas of the video industry will be affected by IP technology.

In this chapter, we will look at the methods used today for delivering video signals to viewers around the world. Then we'll discuss the main technologies used in telecommunications networks. We will also investigate some of the issues surrounding video transport on the Internet. Finally, we will introduce our Video User Checklist, which will be augmented throughout the book. By the end of this chapter, you should be familiar with the common forms of video transport and the common types of telecom networks.

DEFINING IPTV

Problems can occur when new terminology is created and not everyone agrees on the meanings. Case in point: the term *IPTV*. While it is true that all Internet Protocol Television (IPTV) installations send video over IP networks, it is not true that any kind of video sent over an IP network is IPTV. For the latter, the term *Internet video* is much more descriptive.

IPTV is simply a way to deliver traditional broadcast channels to consumers over an IP network in place of terrestrial broadcast, CATV, and satellite services. Even though IP is used, the public Internet actually doesn't play much of a role. In fact, IPTV services are almost exclusively delivered over private IP networks, such as those being constructed by telephone companies in the United States and elsewhere. At the viewer's home, a set-top box is installed to take the incoming IPTV feed and convert it into standard video signals that can be fed to a consumer television.

Some of the main characteristics of IPTV include:

- Continuous streams of professionally produced content (such as a TV broadcast network feed)
- Hundreds of 24×7 channels
- Uniform content format (all channels typically share one compression method and use roughly the same bit rate)

- Delivered over a private network, such as a telco digital subscriber line (DSL)
- Viewed on consumer televisions by way of a set-top box

INTERNET VIDEO

Internet video is used to supply video content to viewers by way of the public Internet. In a typical Internet video installation, service providers set up a website portal that can be reached by anyone with a standard browser. At this site, a list or index of the various pieces of content will be available. Once the user has selected content, it is delivered from servers to the viewer's PC, where media viewer software can be used or where it can be downloaded to another device.

Some of the main characteristics of Internet Video include:

- Discrete content elements, ranging from clips lasting a handful of seconds to full-length movies
- Millions of content offerings
- Widely varying content formats, including dozens of different types of video compression, rights management technologies, and image resolutions
- Delivered over the public Internet
- Viewed on PCs via software, on portable video players, or on televisions by means of network adapters

VIDEO TRANSPORT TECHNOLOGIES

Television was invented for a single purpose: to transport moving pictures from one location to another. The word *television* comes from the Greek word *tele*, meaning "distant," and the Latin verb *visio*, meaning "to see." (The word *video* also comes from this same Latin root.) Hence, *television* means "seeing at a distance." Modern video transport technology is all about sending moving images to a viewer who is far away.

Today, users have many options for video transport, which presents a challenge: How does a user select the best way to transport video

for each application? Many factors are involved in this choice, so there is no one best answer for all users. Let's begin by looking at the many methods used for transporting video signals today.

Broadcast TV

The first IP networking applications in broadcast TV production were video file storage and retrieval, particularly for supporting digital video-editing stations. From there, it was a small step to supporting *live ingest*, which is the process of taking live video content into the digital domain for storage in files on hard disks. Then live-to-air output over IP networks became feasible. This whole process has been driven in part by the continuing spread of high-performance computer workstations that are able to handle video streams in real time. It is actually simpler to configure these workstations with a high-bandwidth networking card (such as Gigabit Ethernet) than it is to equip each station with a video card and audio input and output cards.

Many people first encountered TV in its original form—as a signal broadcast from a central tower, through the air over a dedicated channel, to a television equipped with an antenna. Let's look at some of the key components of a modern broadcast television system (see Figure 1-1).

The master control room (MCR) is the operational hub of a television station, where video content is assembled and made ready for broadcasting to viewers. Video content can come from a variety of sources—a live local source, a broadcast network,¹ a videotape, or from a video server. Table 1-1 gives a small sample of the many functions that a modern television station must perform.

1. Somewhat confusingly, the term *network* has two different meanings for people with broadcast television or data communications backgrounds. In this book, we will try to use the term *broadcast network* whenever we are referring to the distributor of programming that operate in many countries, such as the BBC in the UK, ARD in Germany, or CBS in the United States. When we use the term *network* by itself or with another modifier, we are referring to a data, voice, or other type of telecommunications system.




FIGURE 1-1 Broadcast Television System Diagram

A studio-to-transmitter link (STL) is used whenever the master control room is separated from the actual transmitter tower. The STL carries the television signal directly to the transmitter, normally over a dedicated facility. Microwave radio can be used where there is a direct line of sight available from the studio to the transmitter. Fiber-optic links are more reliable than microwave radio links, but they require a fiber-optic cable to be connected from the studio to the transmitter. These links can be owned by a local telephone company, some other local utility, a municipality, or even by the television station itself.

At the transmitter, the signal is received from the studio and then placed into a specific channel frequency band. For example; in the

TABLE 1-1

Television Station Functions

- Collect video content from a variety of sources, including broadcast network feeds, advertising agencies, local television studios, and syndicated program distributors.
- Prepare the video content for broadcast by editing the material to fit into time constraints and adding local programming and advertising.
- Ensure that the broadcast signal meets all of the performance standards (such as operating frequency and peak radiated power) specified in the station's broadcast license.
- Make sure there is no "dead air," i.e., times when nothing is being broadcast.

United States, Channel 30 occupies the frequencies between 566 and 572 MHz. The modulated signal is amplified to a high power level and fed into the broadcast antenna. The television signal then radiates from the antenna to viewers' locations.

At each viewer's location, a receiving antenna collects the radiated signal and generates a very tiny output signal. This signal is then demodulated and decoded to recover the video and audio signals. These signals are then amplified many times over until they are powerful enough to drive the television set's display and loudspeakers.

Satellite Television

Satellites orbiting the earth are commonly used to receive signals from one earth location and send them to another. Satellites can also be used to broadcast a collection of television signals directly to viewers' homes. Both these applications are in widespread use today.

Satellite transmission of television programs has been used since the mid-1960s to send live programming from one continent to another. Broadcast networks began using satellites to send programming to local television stations and cable TV systems in the mid-1970s. It wasn't until the late 1980s that satellite broadcasting to consumers really began, as exemplified by Sky television in the UK in 1989. This market became known as the direct-to-home (DTH) market, since satellite television service providers were transmitting their programs directly to consumers rather than to local television broadcast stations or cable television systems as in the past. This service is also commonly known as direct broadcast satellite (DBS) service. Let's look at the key components of a typical satellite DTH system (Figure 1-2).

An uplink facility transmits signals from the ground to a satellite, using a high-power signal and a large-diameter dish. The uplink facility gathers video content from a variety of sources, including local television stations in various cities, specialized programmers (such as movie and sports networks), and many others. Because a




FIGURE 1-2 Direct-to-Home (DTH) Satellite Television System Diagram

single uplink facility can create multiple signals to be sent to one or more satellites (by means of a separate dish for each satellite) for rebroadcast over a large area, one facility can serve an entire continent.

Satellites are positioned above the equator at a height of 22,300 miles, which causes them to orbit the earth once every 24 hours and appear to be stationary above a fixed point on the earth's surface. Each satellite is equipped with multiple transponders, each of which receives a signal transmitted by an uplink, amplifies it, and broadcasts the signals back to earth. Current transponders don't do any processing of the uplinked signal (other than amplification and frequency conversion), so a transponder can be used for any type of content, such as voice, video, or data. This permits a DTH broadcaster to change over time the types of services offered without modifying the satellite.

One of the biggest changes in DTH satellite broadcasting was the conversion from analog to digital transmission, which allowed satellite service providers to put from 8 to 20 (or more) television channels on each transponder. Using more channels per transponder increases the number of channels that the DTH broadcaster can offer. One of the first satellite providers to offer exclusively digital satellite service was DirecTV, which began operation in 1994. All of the video signals

are digitized and compressed in a digital satellite system, which helps to simplify the migration to high-definition (HD) content.

Each satellite TV customer must purchase and install a dish antenna. The antenna must normally be installed outdoors with a clear line of sight to the satellite, with no intervening buildings, branches, leaves, or other objects. The dish must be installed at the correct elevation angle (up and down tilt) and azimuth (compass direction) to point directly at the satellite. The LNB (low-noise block) converter in the dish assembly feeds an output signal to the satellite receiver, which performs the functions listed in Table 1-2.

Private satellite networks are also used heavily by television broadcasters, both for gathering content and for distributing it to local broadcasters. When content is being gathered, remote locations take turns in transmitting their content to a satellite, which retransmits the signal to the main studio. (By using only a single transponder, the broadcaster's rental costs for the satellite are minimized.) This method is commonly used for live events in remote locations, such as sports arenas. In many cases this content is scrambled to prevent unauthorized viewing of the unedited video feeds. For distribution, the main studio uplinks a signal to a satellite that rebroadcasts the signal to multiple local television stations. Much of this content is also scrambled. In many cases, content is sent to the local stations

TABLE 1-2

Satellite Receiver Functions

-
- Accept commands from the customer's remote control to select the desired programming.
 - Ensure that the customer has authorization to watch the selected channel (sometimes by means of a special identification card).
 - Receive signals from the LNB, and recover the desired channel. This step may include descrambling the incoming signal.
 - In many cases, the satellite receiver must be connected to a telephone line to permit communication with the DTH service provider's central computers for billing, authorization, and other functions.
-

For digital programming, the following functions must also be performed:

- Demodulate and demultiplex the digital data to select the correct digital stream.
 - Remove any encryption that has been performed on the signal.
 - Decompress the stream to create a video output that is fed to the customer's television set.
-

in advance of when it is needed for broadcast and then simply stored locally and “played out” at the appropriate time. Some corporations also use IP-based networks (often called VSAT, for very small aperture terminal) for data, voice, and video transport applications.

Cable TV

Cable television (CATV) can trace its origins to broadcast television beginning in the 1950s. Originally, the acronym CATV stood for community antenna (or community access) TV, and that was exactly its purpose—to allow viewers to share a common antenna that was able to pull in signals from remote broadcast stations that would be impractical for viewers to receive on their own.

Since those days, CATV has made major strides into new areas, including subscription television (movie channels), pay-per-view, video on demand, and special-interest channels devoted exclusively to sports, weather, or other topics. The video transport mechanism is essentially the same as that used by broadcast television—each video signal is placed onto a specific carrier frequency (TV channel) and transported into each subscriber’s home via a coaxial cable in place of an antenna. Let’s look at a typical CATV system (see Figure 1-3).




FIGURE 1-3 Cable Television System Diagram

A CATV system can serve tens or hundreds of thousands of viewers from a single master headend (MHE), where video content is gathered and made ready to transmit to viewers. Video feeds can come from local TV broadcasters (received at the MHE via an antenna or by means of a dedicated link from the TV station), from satellite feeds, from local programming sources (including educational and government institutions), and from prerecorded content, such as advertisements.

Many CATV providers have started supplying some or most of their programming in digital format, because video streams can be compressed to use network bandwidth more efficiently. This allows the transmission of multiple digital video channels in place of one analog channel and supports the delivery of HDTV. This improvement in efficiency becomes even more important as CATV providers seek to free up bandwidth that can then be used to provide other revenue-generating services, such as voice and data transport.

The techniques used to combine and organize the different digital video signals vary from one CATV provider to another. In Europe, many systems operate in accordance with standards published by the Digital Video Broadcasting (DVB) consortium. Many other techniques are used around the world, including some that are based on IP technology.

The functions of an MHE are very similar to those of a television station's master control room (MCR), with three major exceptions:

- The MHE typically handles dozens or even hundreds of channels, whereas an MCR generally handles one or possibly a handful of channels.
- A portion of the channels processed by the MHE are typically scrambled or encrypted to prevent them from being watched by subscribers who have not paid the appropriate fees. The technology required to manage this content is called a conditional access (CA) system, which controls the decryption and descrambling function for each subscriber's receiver device, commonly known as a set-top box (STB).
- In many cases (such as for pay-per-view movies or cable modem service), communications back from subscribers is permitted.

The MHE must process these return path signals to provide subscriber Internet access or to fulfill orders for pay-per-view movies.

All of the content is modulated onto radio frequency (RF) carriers and combined into a broadband RF output, typically covering a range from 50 to 860 MHz. Some systems operate up to 1 GHz in frequency, with the higher frequencies generally reserved for data and voice services. This broadband signal is then distributed to subscribers using a tree-and-branch type network that typically employs both fiber-optic and coaxial cable, creating a hybrid fiber coax (HFC) network.

From the master headend, a supertrunk is used to distribute the broadband signal to local hubs. Supertrunking may be either analog or digital, but it is almost always fiber-optic. In some CATV systems, processing is done at the hubs to add local programming or to handle cable modem data.

From the hubs, the broadband signal is sent out to the distribution network. The distribution network can be fed directly by coaxial cables, or fiber-optics can be used to cover part of the distance from the hub to the subscriber's home. Customers receive the broadband signal via coaxial cable either directly at their television sets or via STBs.

In many CATV systems, consumers are able to hook a "cable-ready" television set directly to the incoming coaxial cable. In this case, the tuner in the television set is responsible for selecting the correct channel from the incoming broadband signal. Many consumer videotape recorders and DVD recorders are also equipped with cable-ready inputs. In many other cases, customers use a set-top box to receive the broadband signal, particularly when digital transmission is being used. Key functions of a typical STB are given in Table 1-3.

New Technologies

All of the technologies described in the preceding sections currently have hundreds of millions of subscribers around the world. However, some new technologies are gaining popularity. They include

TABLE 1-3

Set-Top Box (STB) Functions

-
- Accept commands from a user's remote control to select the desired programming.
 - Verify that the user has authorization to watch the selected channel.
 - Receive the broadband signal from the coaxial cable network.
 - Tune to the correct RF channel.
 - If the channel is scrambled or encrypted, process the signal to create an unscrambled output.
 - If the channel is digital, demultiplex and decode the digital information into an output suitable for the television display.
 - Generate video and audio signals for display on the customer's television.
-

fiber to the premises (FTTP) and digital subscriber line (DSL). Let's quickly review both of them.

Fiber to the premises (FTTP) technology involves making a direct, fiber-optic connection to each customer's home or business to replace aging, limited-capacity copper wires with new, high-capacity optical fibers. This technology also goes by the names fiber to the home (FTTH) and fiber to the business (FTTB). Recent work on passive optical networks (PONs) has created a new method for delivering video services to the subscriber.

In essence, a PON is an all-optical network with no active components between the service provider and the customer. The network is optical because the path from the service provider to the customer is entirely made up of fiber-optic and optical components. The network is passive because there are no active elements (such as electronics, lasers, optical detectors, or optical amplifiers) between the service provider and the customer.

One key feature of a PON is the use of an optical splitter near the customer premises, which greatly improves the economics of the system. In the case of one popular standard, up to 32 end customers can be served by one fiber from the service provider. A second key feature of a PON network is that the optical fibers and the optical splitter are configured to handle a wide range of laser wavelengths (colors), so new services can easily be deployed. This gives a PON a great deal of flexibility for adapting to future customer needs. Let's look at how a PON system is constructed (see Figure 1-4).

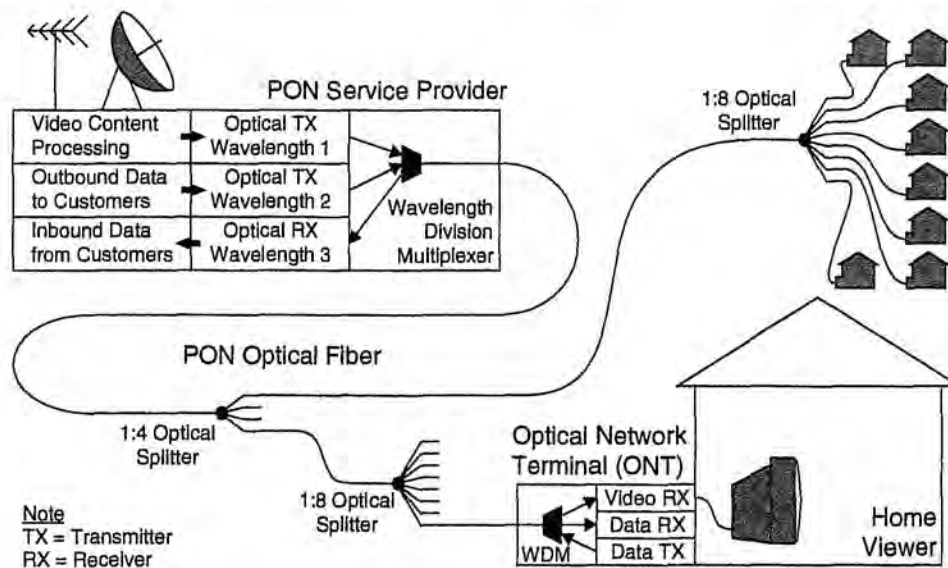


FIGURE 1-4 Diagram of a Fiber to the Premises (FTTP) System Using a Passive Optical Network (PON)

Each customer must have an optical network terminal (ONT; also known as an optical network unit, or ONU) to receive signals from the PON and convert them into electrical signals that can be used by customers' devices. The ONT also converts data being sent by the customer into optical signals that are sent back over the PON to the service provider. In general, the ONT uses normal commercial power supplied by the customer, and it must have a method to power itself (i.e., a battery) to support emergency communication.

The system shown in Figure 1-4 uses three different wavelengths (colors) of light. Two wavelengths go from the service provider to the customers. One wavelength is normally used for a multichannel video signal. A second wavelength carries data from the service provider to all the customers. A third wavelength carries data from the customers back to the service provider. At both the service provider and customer ends of the fiber, specialized filters (called wavelength-division multiplexers, or WDMs) are used to separate the three colors of light.

Video in a PON can be as simple as an optical version of a normal CATV signal that the customer's ONT feeds into a standard cable-ready television set or a set-top box. One big advantage of this

system is that customers may not need a set-top box if they are not interested in scrambled, encrypted, or digital content.

Data in a PON system is multiplexed. From the provider, one stream that contains all of the data for all of the customers is created. Each customer's ONT receives and processes the incoming bit stream and selects the data for that customer. In the reverse direction, only one customer ONT can transmit at a time. For data-intensive customers such as businesses, other wavelengths of light can be used to provide dedicated, high-speed links over PON optical paths.

A major advantage of PON technology is that it can supply very high bit-rate data connections to each customer, with support for aggregate bit rates from the service provider to customers of 622 Mbps and 155 Mbps from the customers back. If all 32 users are running simultaneously at maximum speed, this works out to 19 Mbps that can be transmitted to each subscriber. This amount of bandwidth can easily handle several compressed video signals, so some PON operators simply use this bandwidth to provide IPTV, video-on-demand, and other services.

A major drawback of PON technology is that it requires the service provider to install a fiber-optic connection to each PON customer. This is not a big obstacle in new construction, but the expense of replacing a large installed base of copper cabling with optical fiber and supplying each user with an ONT can be a very large investment for existing customers.

Digital subscriber line (DSL) technology has become popular with service providers and many customers because it offers a way to use existing copper, twisted-pair telephone lines for carrying high-speed data. Many consumers are aware that they can purchase DSL service for Internet access. Some service providers also offer IPTV that is delivered using DSL service. Let's look at how a DSL system can be constructed (see Figure 1-5).

All DSL systems have to make a trade-off between speed and distance. That is, longer distances must operate at lower bit rates, because losses in the cable increase as the bit rate increases. (As technology improves, these limits change, but designers still need to

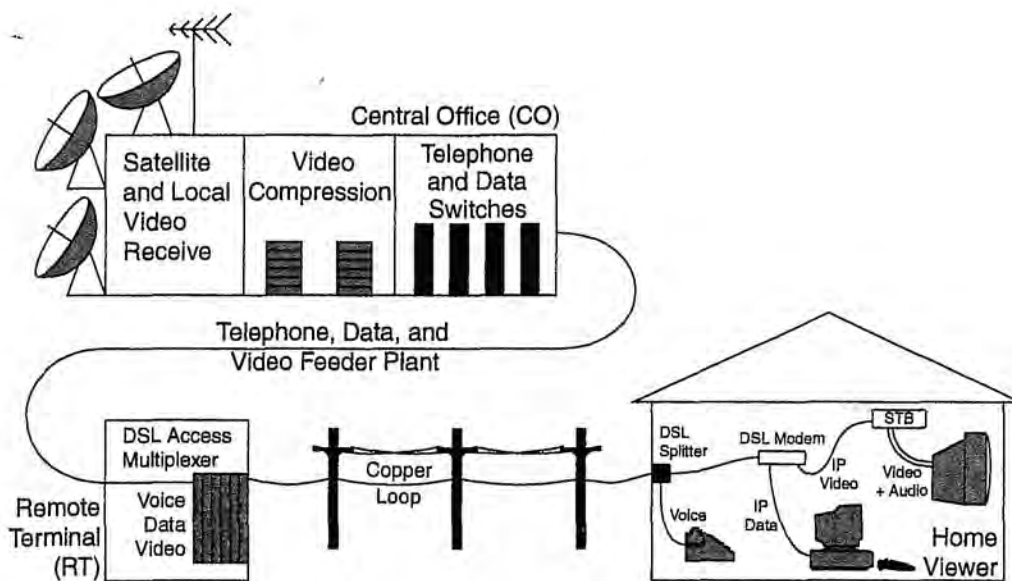


FIGURE 1-5 Digital Subscriber Line (DSL) Diagram

make compromises.) To help keep distances down, a service provider will typically install DSL equipment in remote terminals (RTs), which are located between the provider's main offices and customers.

The central office (CO) is the main hub for services. This is the location of the telephone switch that handles normal telephone traffic. The CO is also usually the place where connections are made to the Internet and where video services are processed prior to distribution.

Customers who are located close to the CO can be fed directly from copper pairs that terminate in the CO. More distant customers are typically served by an RT. Fiber-optics are commonly used to send the signals from the CO to the RT; this is called the *feeder plant*. In the feeder plant, voice, video, and data signals often travel over different transmission equipment. Feeders for video and data services are generally installed alongside existing voice feeders and typically carry much higher-rate data signals.

Each RT provides direct copper connections to each customer in the RT serving area. When DSL service is installed, a digital subscriber line access multiplexer (DSLAM) is also placed into the RT (or in a

CO for directly fed customers). The DSLAM takes video and data signals from the feeder and selects the correct ones for each customer. Then the DSLAM generates the DSL signals and places them onto the pair of copper wires (or local loop) leading to each home.

Every DSL customer must install a DSL modem. The modem receives the DSL signals from the DSLAM and converts them into the proper form for the customer's other devices, such as a PC, a data router, or an IPTV STB. The modem also takes data signals from the customer and transmits them back to the service provider. Table 1-4 lists some of the more common types of DSL services that are available.

One big advantage of a DSL system is that it uses the existing wires that already run to homes and businesses for telephone service. Another advantage is that the DSL circuits are normally designed to fail gracefully, so if a customer loses power or the DSL equipment fails, then normal telephone calls can still be made.

A disadvantage of DSL services for video is that only a few video signals at a time can be sent down a DSL line. Typically, these technologies are restricted to one video content stream per video device, each of which must be equipped with a set-top box (STB). The STB is

TABLE 1-4

Common Types of Digital Subscriber Line (DSL) Services

-
- Asymmetric DSL (ADSL) operates at speeds up to 8 Mbps (millions of bits per second) from the CO to the customer (downstream) and up to 800 kbps (thousands of bits per second) from the customer to the CO (upstream). Two newer technologies offer higher speeds: ADSL2 has up to 12 Mbps downstream and 1 Mbps upstream; ADSL2+ doubles the downstream speed to 24 Mbps.
 - G.Lite DSL (also known as universal DSL) operates at speeds up to 1.5 Mbps downstream and up to 512 kbps upstream.
 - High-speed DSL (HDSL) operates symmetrically (same speed upstream and downstream) at a rate of 1.544 or 2.048 Mbps and is mostly used to provide services to businesses.
 - Very-high-speed DSL (VDSL) operates at bit rates from 4 to 52 Mbps downstream and 1.5 to 2.3 Mbps upstream, although the higher rates can be achieved only on very short links (less than 300 m). Newer standards have increased both upstream and downstream rates, but distances remain short.
 - xDSL collectively refers to the preceding standards as a generic term.
-

Note that the actual bit rates that can be achieved on a DSL circuit are highly variable and depend on many factors, including the length of the subscriber's loop and the amount of noise or interference present on the line.

responsible for receiving the digital video signal from the DSLAM and converting it into the correct output for the user's video display. The STB is responsible for accepting channel-change commands from the viewer and sending them to the DSLAM. Each time a customer wishes to switch the program being viewed, the DSLAM must switch to send out a different video stream.

Dedicated Networks

In the world of television, movie, and video production, special-purpose links are often used to transport video. Table 1-5 lists several advantages of dedicated networks for corporate users.

Of course, all these advantages come at a price—monthly rental fees that usually run between \$1,000 and \$5,000 per month per dedicated TV1 video link for local service in major U.S. cities. These networks are not very flexible—installation fees normally apply, and installation lead times usually exceed a month. However, for high-end users (television stations and broadcast networks), these links are extremely valuable and widely used.

Fiber-optic links form the majority of dedicated video networks for analog and digital applications. Fiber-optic links can run over long distances (over 100 km between signal regenerators) and provide many services (including nonvideo services) on a single fiber.

TABLE 1-5

Advantages of Dedicated Video Networks

- Extremely large bandwidths as needed for uncompressed video signals (such as those described in Chapter 3).
- High-quality transport with extremely low error rates, allowing original content to be transported with essentially no degradation.
- Compatibility with video formats that are used in video production, including feeds from studio-quality cameras and professional editing systems, making it easy for video producers to use these services.
- Privacy, since many of these links operate over dedicated facilities that are extremely difficult for malicious users to access.
- Reliability, because these networks are normally isolated from traffic created by other users.

Fiber-optic links also offer much lower error rates than other technologies. Terminal devices have become very inexpensive over the past decade, which helps keep service provider costs down. Drawbacks include the need to connect fibers to both signal source and destination and the inability of some end devices to share fibers with other types of equipment. Also, most service providers offer these transport services only on a local basis; long-distance video transport is generally done over telecom networks.

TELECOM NETWORKS

Many different types of telecom networks are in use around the world today, and most of them can be used for transporting video signals. In fact, it's not at all unusual for one type of network traffic (e.g., IP) to be transported over another type of network (e.g., SONET) at some point in a long-distance circuit. So it's important to have some understanding about each of the network types and how they can affect video traffic. Let's look at some of the most commonly used networks.

Table 1-6 summarizes much of the data in the following four sections.

PDH

The original digital telecom standards were organized by means of a Plesiochronous Digital Hierarchy (PDH), which means "nearly synchronized." More technically, two circuits are considered plesiochronous if their bit rates fall within a strictly limited tolerance range. By far, the most common PDH signal in North America is a T1 (also known as a DS1), which operates at a rate of 1.544 Mbps. This is enough capacity to carry 24 voice channels, each operating at 64 kbps, plus overhead. In Europe, the corresponding signal is commonly called an E1, which operates at 2.048 Mbps and has the capacity to carry 30 voice channels (also at 64 kbps each) plus overhead. Moving up in speed in North America is a DS3 (or T3), which operates at 44.736 Mbps and carries 28 DS1's, equivalent to 672 voice channels. In Europe, the E3 rate is used, which operates at a speed

TABLE 1-6

Telecom Standards Comparison

Bit Rate	PDH		ISDN		SDH	SONET	Voice Channels
	USA	Europe	USA	Europe			
64 kbps	DS0	E0					1
144 kbps			BRI	BRI			2
1.544 Mbps			PRI				23
1.544 Mbps	T1/DS1						24
2.048 Mbps		E1		PRI			30
34.368 Mbps		E3					480
44.736 Mbps	DS3/T3						672
							Payload Rate
51.840 Mbps						STS-1	50.112 Mbps
155.520 Mbps					STM-1	OC-3	150.336 Mbps
622.080 Mbps					STM-4	OC-12	601.344 Mbps
2.48832 Gbps					STM-16	OC-48	2405.376 Mbps
9.95328 Gbps					STM-64	OC-192	9621.504 Mbps

Commonly Used Video Rates

Videoconferencing: 128 kbps to 768 kbps

Compressed video: 1.5 to 15 Mbps

Uncompressed video: 270 Mbps, 1.485 Gbps

kbps: kilobits (1,000 bits) per second

Mbps: megabits (1,000,000 bits) per second

Gbps: gigabits (1,000,000,000 bits) per second

of 34.368 Mbps and carries 16 E1s, equivalent to 480 voice channels. Higher-speed PDH interfaces are no longer used; they have been replaced by SONET/SDH standards.²

PDH standards are still very important today, even though a lot of new technology has been developed. T1, E1, DS3, and E3 rates are still with us because of the huge installed base of networking equipment with these interfaces. Most service providers offer circuits at these bit rates to customers anywhere in their service areas. Also, these bit rates make sense for a lot of voice and data applications (and even for a few video ones) because they are relatively inexpensive.

² Occasionally, you might see a reference to a circuit called a DS0 in the United States or E0 in Europe. This is simply a single voice channel, which operates at 64 kbps.

ISDN

The Integrated Services Digital Network (ISDN) provides two subscriber interfaces—the basic rate interface (BRI) and the primary rate interface (PRI). A BRI has two B-channels that operate at 64 kbps and one D-channel that operates at 16 kbps, for a total of 144 kbps. The PRI speed depends on its location. In North America, a PRI has 23 B-channels that operate at 64 kbps and one D-channel that operates at 64 kbps, for a total of 1.544 Mbps, including overhead. In Europe, a PRI has 30 B-channels that operate at 64 kbps and one D-channel that operates at 64 kbps, for a total of 2.048 Mbps, including overhead.

ISDN circuits can be dedicated or dial-up circuits. Dedicated circuits are similar to most other network connections—an always-on, fixed-bandwidth network link. Dial-up circuits allow temporary connections to be made from one ISDN device to another by use of special network connection procedures, which are similar to those used in placing a normal voice telephone call.

ISDN lines have some video transport applications. The H.320 videoconferencing standard (see Chapter 10) uses ISDN lines for call setup and video transport. Also, some businesses use PRI lines for data traffic and Internet access, and low-bit-rate IP video signals can flow over these links. Due to relatively high prices and low bandwidth, many ISDN services are being replaced with IP services, particularly DSL.

SONET/SDH

The Synchronous Optical Network (SONET) standard in North America and the Synchronous Digital Hierarchy (SDH) standard in Europe are based on fiber-optic technology. These networks can operate at speeds ranging up to (and beyond) 10 Gbps. What's more, SONET and SDH networks have a common set of bit rates that allow many pieces of equipment to change from one standard to the other.

Because it is easier to understand, let's look at the SDH standard first. The basic SDH building block is an STM-1, which stands for Synchronous Transport Module-1. An STM-1 operates at

155.52 Mbps and can contain a huge variety of different payloads. Mappings, or schemes for constructing payloads, are available for voice signals, data signals (including ATM and IP, see the following section), video signals, and others. STM-1 signals can be combined to form higher-speed interfaces, including STM-4 (622.08 Mbps), STM-16 (2.48832 Gbps), and STM-64 (9.95328 Gbps). Even higher speeds are possible in the future.

Now let's look at the more confusing world of SONET. The basic building block is the STS-1 (Synchronous Transport Signal), which operates at 51.84 Mbps. Because this rate is so close to DS3, it was never really offered as a stand-alone service. Instead, basically everyone agreed to use the STS-3 rate of 155.52 Mbps. (Note that this is exactly the same bit rate as an STM-1 in SDH.) Since these signals are normally provided over optical fiber, the common name for this signal is OC-3, for Optical Carrier 3.

Higher-speed links are also available. OC-12 operates at 622.08 Mbps, OC-48 at 2.48832 Gbps, and OC-192 9.95328 Gbps. Even higher speeds, such as OC-768, will become common in the future. Again, notice that there is a match between some OC signal rates and STM rates. The concept of the STS is important for only one reason—data inside an OC-*x* is organized into frames based on the STS-1. So an OC-12 is basically made up of 12 independent STS-1 signals, each operating at 51.84 Mbps. For higher-speed data that won't fit into a single STS-1, several of them can be concatenated to form a single, high-speed payload envelope. Concatenated signals are denoted by adding a lower-case "c" to the signal name. In order to form an envelope capable of carrying 140 Mbps of data you would need to use an OC-3c.

Most long-distance telecommunication networks operate on SONET/SDH backbones today. This has been caused by the growth of fiber-optic technology and by service providers who require their vendors to adopt these standards. So chances are pretty high that, if you are running a signal over a large geographic area or connecting to the Internet, your data will flow over a SONET/SDH network for at least part of the journey. This is true even if you are using other signal types, such as PDH, ISDN, ATM, or even IP, because there are widely used standards for inserting all of these types of signals into SONET/SDH networks.

ATM

Asynchronous Transfer Mode (ATM) was originally created to be a broadband form of ISDN. ATM was quite popular in the 1990s and beyond, but new deployments of the technology have slowed down recently due to the higher costs as compared to IP.

The core concept of ATM is a "cell," which contains 48 bytes of data and 5 bytes of control and routing information, for a total of 53 bytes in a cell. Terminal devices accept user data, break it up into 48-byte chunks, and then add the 5 bytes of header information. The way that the data is broken up depends on the application; so different ATM Adaptation Layers (AALs) have been defined for different uses, with AAL-1 and AAL-5 being the most popular for video. The resulting cells can be combined with cells from other data sources to form an outgoing signal and passed into the ATM network. After each hop along the network, the header of each cell must be processed to make it ready for the next hop along the route. At the far end of the network, cell headers are removed, cell data is extracted, and the data is reassembled into the original format.

Because ATM is a connection-oriented system, before any data can flow across the network, a connection must be established between the source and the destination for the data. This connection establishes a specific route for the data called a virtual circuit (VC). Note that the process of establishing a VC requires stacks of software in the terminal devices and in each piece of network equipment along the route of the signal. Once the VC is fully established, data can flow. Software is also required to handle cell errors and to reestablish connections when failures occur or when a new VC needs to be set up.

Some service providers still use ATM within their own networks, even for non-ATM data (some IP traffic is carried over ATM facilities within carriers). The biggest benefits of ATM for a carrier are the abilities to control the amount of bandwidth used by each customer, to ensure that data from different customers are kept separate, and to ensure that adequate bandwidth exists across the entire data path from source to destination, coupled with the ability to make sure that

certain types of time-sensitive traffic (such as video) have a higher priority over other types of traffic. These properties make it very easy for end devices to send and receive video smoothly and reliably.

The biggest drawback of ATM for most users is cost. First, the user must purchase an ATM multiplexer to format the data into cells. ATM devices can be quite expensive in comparison to IP devices, due, in part, to the large amounts of software required to set up and manage connections. Each user must lease an ATM access link from a carrier at each end of the data circuit. Furthermore, many carriers charge for actual network usage on the ATM links rather than the flat-rate services that are common on IP links. All these factors combine to make ATM transport typically more expensive and increasingly less popular than other services, such as IP.

IP

Since Chapter 5 is devoted to IP technology, this will be a brief introduction. IP is a standard method for formatting and addressing data packets in a large, multifunction network, such as the Internet. A packet is a variable-length unit of information (a collection of bytes) in a well-defined format that can be sent across an IP network. Typically, a message such as e-mail or a video signal will be broken up into several IP packets. Each packet must carry a destination address so that the network knows where to deliver the packet. IP can be used on many different network technologies, such as Ethernet LANs and wireless Wi-Fi links.

To help illustrate the functions of IP, an analogy might be useful. For folks who live in the United States, use of the U.S. Postal Service requires the sender to put a destination address that includes a 5-digit Zip Code on each piece of mail that is sent out. There are rules for assigning IP addresses, and their format is precisely defined. There are many excellent references on this, so we won't go into detail here. Certain IP addresses are reserved for special purposes, such as multicasting (see Chapter 9). It is the same with Zip Codes—the first two digits indicate the state (or part of a big state), the third is the region within the state, and the last two indicate the local post office.

IP is connectionless, which means that each packet can be sent from source to destination along any path that happens to be available. Each packet is handled individually; that is, the network is not required always to select the same path to send packets to the same destination. This can sometimes result with the packets being received in a different order than they were sent. In the U.S. mail, each letter is handled individually. At each sorting location, the destination Zip Code is examined on each letter. The sorter's function is to choose any available method (airplane, truck, etc.) to move the letter along toward its destination. If a particular method is full to capacity, another method will be chosen. This is why two letters mailed on the same day from the same location to the same recipient may not arrive on the same day. IP packets are variable in size. The length of each packet is included in the header information. Of course, mail messages can be of different sizes as well.

IP has some functionality for controlling which packets get priority when congestion occurs. This allows certain packets to be flagged as having high priority and others as low priority. In the U.S. mail, letters and parcels can also be assigned a priority, such as First Class, Bulk Rate, and Express Mail.

Video transport applications on IP networks range all the way from low-resolution, low-frame-rate applications like webcams to high-definition television and medical images. It is paramount that users, technologists, and service providers understand how video works on an IP network. That is a topic we will explore in great depth in the rest of this book.

THE INTERNET

The Internet is a worldwide network that provides IP-based communication services for a wide range of users and applications. One of the most familiar uses of the Internet is the World Wide Web, which can be accessed through web browsers essentially everywhere on the planet. The Internet also supports e-mail, file transfer, and many other applications. In this book, we will use Internet with a capital "I" when we are referring to this enormous, public network.

Many users would like to be able to transport high-quality video signals over the Internet. This is certainly possible and works well for prerecorded content and live content if perfect broadcast quality is not required. For the latter, private networks are typically used. Use of the Internet for video transport has both advantages and disadvantages, as detailed next.

Advantages

Ubiquitous availability—Connections to the Internet are available throughout the world (even Antarctica), and today there are hundreds of millions of regular Internet users. The Internet has many more connections than even the largest private or government network in existence.

Low cost—The cost of receiving data over the Internet is transparent to many users. A monthly fee for access through their Internet service provider (ISP) is all that most users pay. Contrast this with the cost that would be incurred to construct a private network with dedicated facilities to each customer in a worldwide video IP network.

Disadvantages

Unpredictable performance—Because the Internet's function is dependent on the interaction between so many different networks, congestion or delay can occur in many places. Also, there is no mechanism for a user to request a higher-priority level for certain types of traffic, as is often done for video traffic on private networks. These factors combine to make it very difficult to guarantee a required level of performance for any Internet connection.

No multicasting—Multicasting is a powerful tool for sending the same content to multiple destinations simultaneously, as will be described in Chapter 9. Unfortunately, the Internet is not enabled for multicasting, so each video stream needs to be sent individually from each source to each destination.

Overall, it is possible to transport video content over the Internet, and we will discuss several examples in this book. However, users

must be prepared to deal with the shortcomings of the Internet and understand that many video applications will work best on specialized or private IP networks.

REVIEW AND VIDEO USER CHECKLIST

In this chapter, we began with a definition of Internet TV and IPTV, followed by descriptions of the three most popular methods for delivering broadcast television to consumers: terrestrial broadcasting, direct-to-home (DTH) satellite broadcasting, and cable television (CATV). We then examined two new technologies that can be used for consumer delivery: Passive optical networks (PONs) and digital subscriber line (DSL). We also glanced at dedicated video networks for delivering video to broadcasters. Next, we looked at some popular telecommunication networks, including PDH, SONET/SDH, and ATM. We concluded with a brief look at the basics of IP networks and discussed using the Internet for video transport.

Throughout this book, we will be covering many different aspects of video transport over IP networks. These issues affect the choices that users need to make in order to achieve success with IP video networking. To help guide the discussion, we will be creating a checklist of information that users should gather before beginning a video transport project. Following are some of the topics that will make up our checklist.

Chapter 1 Checklist Update

- ☐ *Source of Content:* Who will be supplying the content? Who owns the content? Are there any restrictions on use of the content? Can only certain users view the content? Will the content need to be protected against copying?
- ☐ *Type of Content:* What types of scenes are included in the content? How much detail is there? How much motion is present? Does the detail level of the content need to be preserved for the viewer?
- ☐ *Content Technical Requirements:* Does the content come from film, videotape, or a live camera? Is there a synchronized audio track? Is there more than one audio track (second language,

commentary)? Is there any data that must be included with the content, such as closed captioning, V-chip data, or program descriptions? Are there any limits mandated by the content owner on the amount or type of compression that can be used?

- ☐ *System Funding:* How will the content be paid for? How will the network usage be paid for? Will payments be required from each viewer? Will advertising be used?
- ☐ *Viewer Profile:* How many users will be viewing the content? Where will they be located? What equipment will they use to view the content?
- ☐ *Network Capabilities:* What bit rates will the network support? Will the network support multicasting? What security features does the network offer? Who owns the network?
- ☐ *Performance Requirements:* How much delay will be acceptable in this application? What video and audio quality levels will users accept? Will all users get the content at the same time? Will users be allowed to pause, rewind, and fast-forward the content? What is the financial impact of a service interruption?

A more detailed list appears at the back of this book in Appendix B. In each chapter, we will be adding to this checklist as our discussion covers more topics.

13

IPTV—DELIVERING TELEVISION TO CONSUMERS

IPTV is a method for delivering traditional, linear television programming to consumers over a private IP network. It can be considered a direct substitute for broadcast, satellite, or CATV systems, because it accomplishes the same purpose: delivering hundreds of channels of continuous programming that are displayed on a television by way of an STB. Contrast this with Internet video, which consists of millions of video clips that are viewed by means of a PC display and delivered over a public network. (Internet Video will be covered in Chapter 15.)

The distinctions between IPTV and Internet video are not merely semantic; there are significant differences in the technologies used to provide these services. In particular, the delivery networks are very different, with IPTV hardware and software systems typically custom built for video delivery and Internet video simply leveraging existing infrastructure. Business models differ greatly: IPTV is almost always a paid subscription service, where many Internet

video services are provided for free to the user. Viewer expectations are also divergent—it's not an issue if a few seconds of buffering are required before an Internet video clip begins to play, but no IPTV system would survive long if it introduced 5–10 seconds of “dead air” between each program or every time the channel was changed.

In this chapter, we will start out by defining more completely what IPTV is and then look at two IPTV applications. We'll explore the reasons for using IPTV. Then we will examine some of the key technologies that support IPTV. We'll conclude with an example of an actual IPTV delivery system.

Service providers who wish to deliver multiple consumer services over a single network often choose IP technology because a single platform can provide voice and high-speed data access in addition to IPTV. In a typical system, a private, high-speed IP network is used to deliver video programming continuously to hundreds or thousands of viewers simultaneously. Figure 13-1 shows a typical IPTV network.




FIGURE 13-1 Typical IPTV Network

CHARACTERISTICS OF IPTV

Several key characteristics define the ways in which IPTV is different from other video applications that use IP networks:

Continuous content streams: IPTV is designed to send streams of video programming to each viewer. These streams are continuous; that is, each viewer can select the stream he or she wants to view but must join the stream in progress. This process is functionally identical to the programming delivered by local broadcasters, CATV companies, and satellite providers, where the viewer is able to select the channel to be viewed but not the content of the channels. Of course, many IPTV providers also offer video-on-demand (VOD) and interactive services using their IPTV networks, but these services can also be delivered in many other ways.

Multiple channels: The primary content delivered over an IPTV network is produced by a range of broadcast networks and delivered simultaneously to a large number of viewers. Viewers typically choose which channel they want to watch on their television by interacting with the IPTV set-top box (STB). This can be done simply by entering the desired channel number on a remote control keypad or by making a selection from an Electronic Program Guide (EPG).

Uniform content format: Most IPTV systems use only one video encoding format (or possibly two) for each type of content. The choices typically range from MPEG-2 or MPEG-4 to VC-1. IPTV providers will typically choose one compression format and one bit rate for all SD video signals and another combination for HD signals. This greatly simplifies the overall management of the IPTV system, allowing for a uniform system design and easing the burden on technicians maintaining the system. This also simplifies the STB design by eliminating the need to support multiple video decompression engines. Of course, this approach requires that any content not in the correct format be converted when it arrives at the IPTV system.

Private delivery network: In order to deliver continuous channels of content to thousands of viewers in a repeatable manner, an IPTV network must be carefully provisioned and controlled. This task is very daunting on a private network, where all of

the video content and other network traffic can be controlled. This task would be impossible on the public Internet.

Viewed on consumer televisions via set-top boxes: The role of an STB is extremely important for an IPTV network. At a minimum, it must receive an incoming IP video stream, reassemble the data packets in the correct order, decode the video signal, and produce an output that can be fed to a television (or projector) for display. The STB normally serves as the terminus for the IPTV network, so it must be able to receive commands from the user's remote control and send them into the network for action. It may also need some built-in intelligence to be able to generate text or other graphics to communicate with the viewer, for functions like the EPG.

APPLICATIONS

Many readers of this book are aware of the rich variety of technologies that can be used for video delivery into the home. With this in mind, let's look at two different types of services that IP networks are currently being used to deliver: entertainment to the home and hotel/motel applications.

Entertainment

Some very large IPTV networks have been installed to provide home video services, similar to those provided by traditional CATV and DTH satellite providers. In all these systems, individual households are supplied with a variety of broadcast television channels in exchange for a monthly subscription fee. These fees pay for three basic costs of the service provider: the cost of the programming, the initial capital cost of installing the video delivery system, and the ongoing costs of performing maintenance on the delivery system and providing customer service.

Typically, entertainment systems provide service for a particular geographic area, such as a city, a town, or a neighborhood (or, in the case of satellite services, a country). In many cases a license or franchise from the local government is required to operate the

system; this may require an annual fee to be paid to the government (called a *franchise fee*). Exclusive rights for a certain territory are sometimes given to the service provider in exchange for certain concessions, such as providing video services for governmental or educational use. However, just because a service provider has an exclusive franchise for an area does not mean that the provider is guaranteed to capture all of the television viewers in that area.

In any video delivery system, a number of potential subscribers will choose not to receive video programming and become actual subscribers. The ratio of actual subscribers to potential subscribers, called the *take rate*, can be influenced by a number of factors. These factors include the availability of competing video services (such as CATV and satellite), the take rates of those other services, the amount of local programming available, and the intensity of the marketing program for the service, among many other factors.

Normally, a service provider does not install all of the equipment needed to service every single residence in a given area, because the take rate will begin at a low percentage. Instead, distribution equipment is added as needed to handle new subscribers as they sign up. However, a large amount of common equipment (for content gathering and system operation, as described later in this chapter) must normally be installed to equip an area for video delivery. Because of this, the system economic breakeven point might not be reached unless the take rate reaches 20–30% or possibly even more. Overall, it is very difficult to design a profitable business plan for a new video-only network unless one out of every five households that can subscribe to the service actually does (a 20% take rate).

IP technology offers a way for these service providers to capture more than just video subscribers: the ability to offer “triple-play” services. For example, many CATV companies are adding IP capabilities to their networks to provide voice over IP and data services. On a properly designed pure IP network, video services can be offered, data services such as Internet access can easily be supported, and voice services can be provided using voice-over-IP technology. With a triple-play network, a service provider can more easily achieve the take rate needed to reach economic breakeven.

Hotel/Motel Video Services

IP video technology can also be applied to delivering in-room entertainment video to hotel guests. This technology can be used to provide all of the video services to each room or just the premium movie channels offered by many hotels. Because the video display in a hotel room is normally a television set, an IP-enabled STB is required in each room to receive the incoming signals and convert them.

Typically, hotel systems are configured to use a video server in each building. The server outputs are connected to a data transport system that delivers the IP streams to each room in the hotel. This can be accomplished via existing telephone wiring if DSL technology is used. IP traffic can also be sent over coaxial cable normally used for distributing television signals in a hotel through the use of IP cable modem technology.

One advantage of IPTV in a hotel setting is due to the growing demand for in-room Internet access. With careful provisioning, it is completely possible to use the same in-hotel DSL circuits to deliver both data and video services to each room. With many hotels upgrading their television systems to offer digital TV and HD video services, IPTV becomes an attractive option. In-room telephones can also be converted to voice over IP, but that trend is not yet strong.

BASIC CONCEPTS

Video delivery to the home has created a variety of new terminology and business models that need to be understood in order to make sense of the reams of articles and company literature that have been produced on this topic. In the following sections, we'll look at some of these new terms and how industry insiders are using them. Along the way, we'll also discuss some of the market forces that are pushing service providers to seek out and apply innovative new technologies, such as video over IP.

Why Use IP for Video to the Home?

The hardware and software systems that deliver terrestrial broadcast, DTH satellite, and CATV services are the products of hundreds

of thousands of hours of development and refinement, and they provide services that are very popular with the home viewing audience. In the face of all this proven technology, the logical question becomes: Why introduce a new technology, such as IP video transport, for video delivery to the home?

The answer is not a single compelling reason but, instead, a combination of factors that can influence this decision. Briefly, these factors are local loop technology, user control, convergence, and flexibility. Let's look at how IP technology can impact each of these factors.

Local telephone technology has been evolving rapidly over the past decade. With the emergence of digital subscriber line (DSL) technology, service providers are now able to deliver thousands of kilobits per second of data to each customer's home. This capacity far exceeds what was possible even a decade ago, when speeds were typically limited to a basic-rate ISDN line operating at 144 kbps or a simple dial-up modem circuit operating at 56 kbps. With the added capacity of DSL, the logical question for providers has become: What other services besides telephony and Internet access can we provide to customers that they will be willing to pay for? The answer, in many cases, is video services. To accomplish this, IPTV has often proven to be the best choice for providing video over DSL circuits. By adding IPTV, service providers can earn additional revenue to help pay for DSL speed upgrades.

User control is another powerful trend in the video delivery business. By enabling users to choose what they want to watch at any time, service providers have been able simultaneously to increase user satisfaction and increase the fees that customers pay. For example, personal video recorders (PVRs, which have been popularized by suppliers such as TiVo) allow users to record television broadcasts automatically on a computer hard disk for later, on-demand playback (sometimes called *time shifting*). This service is not free: Users have to pay a significant sum to purchase the PVR device and then many pay an additional monthly service fee to receive the electronic program guide. Another example is a service being offered by cable television companies in the United States called HBO on Demand, which allows users to select programs from a

list of shows and movies for viewing at any time, in exchange for a substantial monthly fee. This service has become popular with customers in those cable systems that offer it. Clearly, users are demanding, and willing to pay for, programming they can control to meet their viewing habits. IPTV is a great technology for providing this control capability, because every viewer can have a customized stream from the provider.

Another big market driver is convergence. Unfortunately, the term *convergence* got rather a bad reputation during the "dot-com" bubble of the late 1990s. However, the concept is still with us, and it is a powerful force in the video delivery business. By *convergence*, we mean the use of a single network to deliver multiple services to consumers. The term *triple play* has come to mean the simultaneous offering of voice, video, and data services, all billed to the subscriber in a single, convenient monthly package. By offering converged triple-play services, many providers are convinced that they can increase customer retention and reduce the costs of their customer service operations.

As the pace of technical innovation has picked up, so has the need for service providers to add flexibility to their service offerings. From a purely economic standpoint, providers strongly prefer installing technologies that will still be usable in 10–20 years. (Without that kind of life cycle, it is hard to justify the significant up-front capital costs of the network while still being able to offer services at prices that customers will pay.) By making sure that their networks are flexible, service providers can adapt to changing market demands while minimizing the cost of adding these new services. IP networks have certainly proved to be one of the most flexible technologies ever created and should remain that way due to the relentless march of technology in the computer industry. IPTV takes full advantage of this flexibility, making it easy for service providers to achieve a low-cost, multipurpose service delivery network.

Faced with the foregoing market forces, many service providers have decided to take a hard look at using IPTV as their basic technology underlying a variety of other service offerings. Not all video providers will adopt IPTV technology, but many will give it serious consideration.

Video-on-Demand

Video-on-demand (VOD) has a great deal of appeal to many consumers. The luxury of being able to sit down in front of the television and select a program or a movie and begin watching immediately is very appealing, and millions of consumers have installed PVRs for exactly this purpose. In response to these customer trends, IPTV providers need to ensure that their systems have a means to support VOD as a current or future offering.

The basic concept of VOD is based on video programming that is stored and then delivered to a viewer when it is requested. This storage can take the form of a centralized server equipped to send programming simultaneously to hundreds of viewers, or it can take the form of more distributed storage throughout the network. At the limit, individual storage devices for each viewer can be located in individual STBs.

Various forms of VOD have been tried over the years, and most of them still exist in one form or another. Table 13-1 lists the most popular types of VOD services.

One of the big controversies surrounding PVR service is the role of advertising in recorded content. Advertisers have two main concerns:

- *Ad skipping*, where viewers fast-forward through ads. This capability is often listed as the motivation for many consumer PVR purchases.
- *Ad timeliness*, where viewers watch programs at times far removed from their original broadcast date. This is a big concern for advertisers who have their ad campaigns targeted for specific time windows, such as promotional ads for a movie being released to theaters the following day.

Service providers have a limited amount of control over content that has been recorded by viewers on their own devices for later playback. They have only a slight bit more control over PVRs embedded in an STB supplied by the service provider—at least they can ensure that the DRM function is working to protect any

TABLE 13-1

Types of Video-on-Demand Service

Type	Description
True video-on-demand (VOD)	This is the purest form of VOD, where each viewer receives an individual video stream over which he or she has complete control. Viewers are allowed to start, stop, pause, rewind, and fast-forward the content. Viewers typically pay a fee for each title viewed; the charges are either debited from a prepaid account or included on a monthly bill.
Near video-on-demand (NVOD)	Similar to true VOD, but without the individual video stream control capabilities. One common form of NVOD is sometimes called <i>staggercasting</i> , in which multiple copies of a program are played starting at five-minute intervals, thereby limiting any individual viewer to no more than a five-minute wait before his or her program begins to play.
Subscription video-on-demand (SVOD)	Same delivery technology and viewer control as VOD, but with a different payment system. In SVOD, subscribers pay a fixed monthly fee for unlimited access to a library of titles. In many systems, the library is updated monthly.
Free video-on-demand (FVOD)	A variation on VOD where payment is eliminated. In most systems, this content is restricted to long-form advertisements, how-to guides and other low-cost content.
Everything on demand (EOD)	For some technology visionaries, this is the ultimate form of video delivery system, where all programming is available to all viewers at all times.
Personal video recorders (PVRs)	These devices take incoming video programming, compress it, and record it to a hard disk that is typically located either in an STB or a stand-alone device. Viewers then control the PVR to play back content, including pause, fast-forward, and rewind capabilities. With this capability, also called <i>timeshifting</i> , viewers normally program the PVRs to record specific programs at specific times. One of the pioneers of this technology is a company called TiVo.
Network personal video recorders (NPVRs)	Offers functionality similar to PVRs, but recording is performed inside the service provider's network rather than in the viewer's location. Some content owners contend that this technology is so similar in capability to true VOD that it needs to be licensed as such.
Pay per view (PPV)	This precursor technology to VOD is used primarily to deliver live paid programming, such as concerts and sporting events.

copyrighted content while it is on disk. Providers actually have the potential to influence viewers who use a networked PVR, where the video recordings are actually stored on the service providers' own video servers.

Network PVRs have exciting potential to make advertisers much happier than other PVR technologies. Why? Well, consider what happens in a normal PVR scenario with an advertisement. The machine faithfully records the commercials along with the program content and gives the user the ability to fast-forward through any parts of the program or advertisements at their whim. For example, say the viewer recorded a program on December 20 and decides to watch the program on December 29. As you might imagine, the program contained a number of ads that pertained to special last-minute shopping opportunities for Christmas. Unfortunately, when the viewer watches the program, the sales are over and the ads are completely worthless to both the viewer and the advertiser. Now consider the same scenario with a networked PVR and some advanced technology in the server. With this technology, the service provider is able to replace the commercials in the original program with ones that are timely and relevant whenever the viewer watches the content. In our example, the ads might be for something great to do on New Year's Eve, which the viewer might actually be willing to watch and for which an advertiser might be willing to pay.

All that's needed to make this a reality is some pretty serious software inside the VOD server and some kind of legal framework to govern the "bumping" of one commercial by another. The industry isn't quite there yet, but this technology is certain to be available in the not-too-distant future.

Interactivity

Interactivity is like beauty—very much in the eye of the beholder. Because of the many different meanings that can be assigned to this word, let's examine three different ways in which viewers can interact with their televisions. We'll begin with basic interactive functions, then look at "VCR-like" controls, and end with content-oriented interactivity. Each of these terms will be defined in this section.

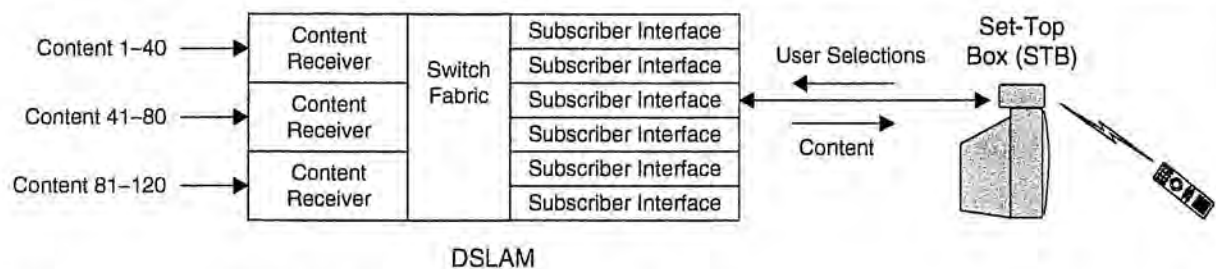


FIGURE 13-2 Example of a DSL Video Delivery System

The most basic way that viewers interact with a television (*basic interactivity*) is to turn it on and off and to select a programming channel to view. In a typical broadcast, DTH satellite, or CATV system, this is a very simple function to implement and can be completely controlled by equipment in the viewer's location, such as a TV tuner or a set-top box. With the current generation of video over DSL, there is usually only enough network bandwidth to deliver one video signal per television. This forces the channel-switching function to move out of the viewer's location and into the provider's facility. The equipment required to implement this for a DSL system is a bit more complex than for a CATV system, as shown in Figure 13-2.

In this example, the viewer interacts with his or her digital STB. Whenever he or she wants to select a new channel, the viewer presses the appropriate command on his or her remote control. The STB takes this command and relays it to the digital subscriber line access multiplexer (DSLAM). As soon as the DSLAM receives the command, it locates the new content and begins sending it to the viewer. This is a relatively simple process for broadcast channels, where the DSLAM is already receiving the content. The switch inside the DSLAM is reconfigured to stop sending one video channel to the viewer and to start sending another, typically by way of an IGMP "Leave" command followed by a "Join" command (see Chapter 9). Although this takes a little bit of time to accomplish, it does not take noticeably more time than switching channels on a digital satellite system.

The next-higher level of interactivity, *VCR-like*, occurs when viewers are given the same set of controls they would have if they were operating their own personal video cassette recorder—that is, start, stop, pause, rewind, and fast-forward. Implementing these functions

puts a much higher burden on the video delivery system, because each viewer needs to be supplied with a video stream that is unique to him or her and completely under his or her control. There are a number of ways to achieve this functionality, as we discussed in the previous section, on VOD. One thing that must be recognized is that the amount of bandwidth needed to provide VOD increases as the number of viewers increases. That is, as more subscribers are added to a DSLAM, chances are more video bandwidth will be needed to feed the DSLAM.

Content-oriented interactivity occurs when users are given the ability to respond directly to the content itself. One example of this would be responding to an advertisement by requesting more information or possibly even placing an order for a product directly through a web interface on the television (a dream of many infomercial vendors). Another example would be influencing the outcome of a video program, such as voting for the winner of a talent contest during a live broadcast and then seeing the results of the voting during the broadcast. Note that for live content, the servers used for VOD do not come into play, because the content must be distributed from the source directly to viewers without being stored on servers. Instead, a mechanism must be provided to transmit the viewers' commands or choices back to the origin of the broadcast. In the case of prerecorded content with different program branches selected by the viewer (similar to DVDs with alternate endings), the VOD server can simply treat each branch of the content as a separate piece of content that can be selected by the viewer, and the interactivity flows would terminate at the server.

The impact of these different levels of interactivity is presented in Figure 13-3, which shows the data flows that occur in an IP-video-over-DSL system under each of the preceding scenarios. For basic interactivity, the data flow starts at the user's remote control, passes through the set-top box, and then moves to the DSLAM, where it terminates. For VCR-like user control, the data starts out the same way but ends up passing through the DSLAM to video servers located further upstream. With content interaction, the viewer responses are sent even further upstream, to the content provider. It is possible that all these different data flows will need to be accommodated in a single video delivery system, because individual subscribers may choose different levels of interactivity.

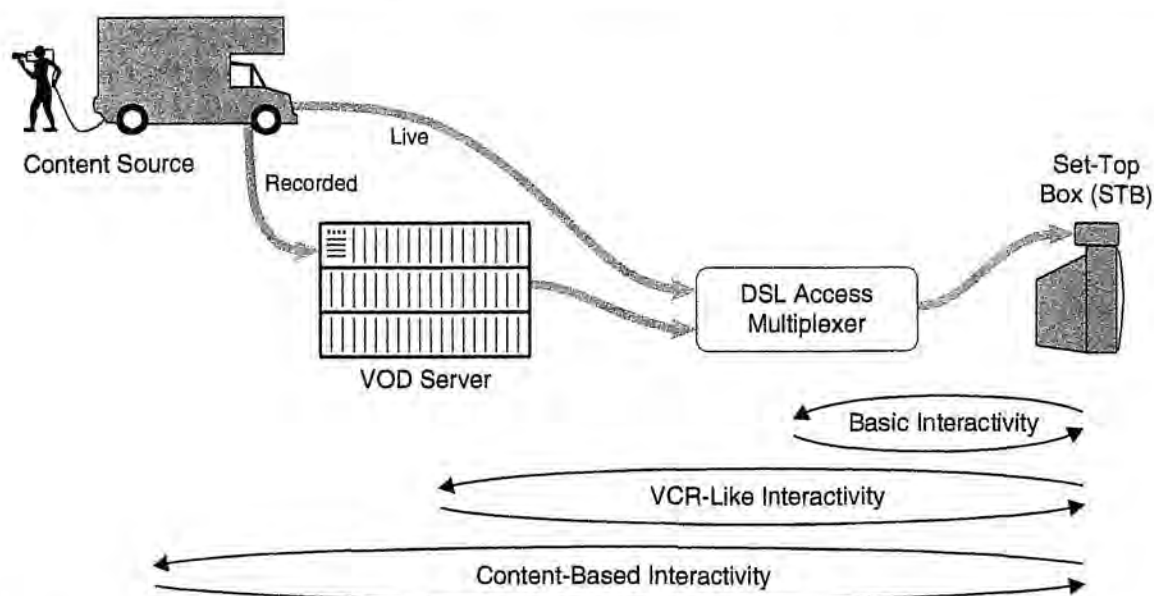


FIGURE 13-3 Interactivity Data Flows

Operations Support Systems

Profitable delivery of video services to the home requires much more than a high-reliability hardware platform. A great deal of software is also required for managing a huge array of tasks, from informing customers about the programming available on the different broadcast channels to capturing the data needed to bill subscribers for the services they have purchased. Collectively, these software systems are known as *Operations Support Systems (OSS)*, and they can take many forms. Table 13-2 lists many of the key functions provided by IPTV OSS systems.

OSS systems can be a major portion of the investment for an IPTV service provider, in terms of both time and money. Because of the wide variety of functions that can be provided, it is likely that software will need to be purchased from multiple vendors to implement the full range of functions selected by the provider. Integrating these systems can take months, and much of this work needs to be completed before large numbers of subscribers can be serviced. Furthermore, these expenses are essentially fixed costs whether the service attracts 1000 or 100,000 subscribers. As such, the costs of installing OSS systems need to be carefully considered in a service provider's business plan,

TABLE 13-2**Key OSS Functions for IPTV Systems**

- An *electronic program guide* gives viewers the programming schedule for each broadcast channel and available VOD titles. This guide can be either a simple broadcast channel that scrolls through all the program choices or an interactive program guide that allows the user to navigate through listings, possibly to look at programming scheduled in the future. Many IPTV providers use outside companies (such as *TV Guide* in the United States) to supply program guide data.
- An *order entry* system is required when customers are permitted to purchase content through the IPTV system. This system needs to be able to verify that the customer's account information is valid and that the order can be filled as requested. This system needs to connect to the subscriber billing system and may also need to connect to external services for credit card validation, etc. A secure (encrypted) link is required to customer homes when personal data is being collected.
- *Online content access* (e-mail, web surfing) is provided by some IPTV systems, allowing viewers to use their television sets for functions that might otherwise require a PC.
- A *subscriber management and billing system* maintains crucial data about each subscriber, including contact and billing details, account status, and equipment identifiers. Many of the other OSS components will refer to this data to obtain information needed to authorize new services, deliver video streams, and so on. Customer service personnel will likely also be heavy users of this system because they support subscribers who move or upgrade or downgrade their services; add or remove televisions; and require repair services.

recognizing that in the early stages of deployment, these costs can exceed the costs of the system hardware for low subscriber counts. Also, the ongoing costs of maintaining the database should not be overlooked when developing a business model for an IPTV system.

IPTV DELIVERY NETWORKS

A wide variety of networks are used in delivering IPTV. There are the basic delivery mechanisms, such as DSL, passive optical networks (PONs), and traditional cable TV lines. In addition, traditional high-speed data networks connecting offices and homes can be used for applications with widely dispersed viewers. We'll take a closer look at all of these network technologies in the following sections.

Digital Subscriber Line (DSL)

Current-generation DSL technologies, such as asymmetric digital subscriber line (ADSL), provide relatively limited amounts of




FIGURE 13-4 DSL Technology Bandwidth Comparison

bandwidth from the service provider to the consumer and even more restricted links from the consumer back to the provider. As Figure 13-4 shows, data rates of all DSL technologies decrease with increasing circuit lengths. Newer technologies, such as ADSL2 and ADSL2+, are being used for most new installations, because they offer more bandwidth in the forward path.

VDSL (very high-speed digital subscriber line) technology supports significantly more bandwidth on each subscriber line. Accordingly, more video channels can be transmitted to each VDSL subscriber, with three or four simultaneous videos possible. A few high-definition video signals could also be transmitted. One drawback to VDSL is that the range of operational distances is less than that for ADSL, so subscribers need to be closer to the service provider facilities. Also note that the speed of DSL services varies with distance, so good planning for varying data rates is essential.

H.264 technology has become critical for delivering HD IPTV services, because signal compression into the range of 4–10 Mbps is required for most DSL technologies. Good quality MPEG-2 compression usually doesn't get below 12–15 Mbps, which exceeds the

bandwidth of many DSL circuits. By using bit rates near the lower end of H.264's range, it becomes possible to have more than one HD feed into each home, which could be used to support two or more HD STBs simultaneously.

Passive Optical Networks (PONs)

Although DSL-based IPTV systems represent the majority of the new circuits being installed today in North America and elsewhere, IPTV can also be implemented over PON systems. Since we talked about PON technology in Chapter 1, we won't repeat that information here.

IPTV technology can be used to distribute one or more video signals to a home via the data portion of a PON. When used in this manner, the IPTV system will look very much like the IPTV system we discussed for DSL: Individual channels will be requested by the viewer at each STB, and the service provider's IP network will deliver the requested video channels. This system merely replaces the DSL link with a PON data link.

Another alternative is a hybrid system. In this system, broadcast channels are delivered by one wavelength of light in the PON system, and the VOD and other user-specific video channels are delivered via the data transport paths. This scheme requires an STB that can accept both broadband video inputs and IPTV inputs. One advantage of this type of a hybrid system is that the IPTV equipment at the central location could be scaled back in capacity, since typically only 5 percent of viewers watch user-specific services at any given time, although this proportion might increase with greater availability of subscription VOD.

IPTV Over Cable

Traditional cable television (CATV) networks have excelled in the delivery of hundreds of channels of video content simultaneously to thousands of users. Each user can select any one of the broadcast

channels available, simply by tuning a television or STB to the correct channel. These systems are easy to scale; adding new customers simply requires splitting and amplifying the signal to reach more households. In the past, interactivity has been extremely limited or not available at all, since all the content was sent in one direction only—to the viewer.

Increasingly, CATV providers have begun to look at how to provide more advanced video delivery systems and ones that allow them to offer triple-play video, voice, and data services. IP technology is a natural platform for converging these different services. Some CATV providers are considering using IPTV over a standard CATV plant.

To make this operate, it is necessary to transmit IP data packets over the cable network. To do this, special digital modulation schemes with acronyms like QAM (quadrature amplitude modulation), OFDM (orthogonal frequency-division modulation), and VSB (vestigial sideband) are used. These technologies give the CATV significant data-handling capacity—on the order of 35–40 Mbps in place of every 6-MHz analog video channel. This is quite attractive; if an average SD digital video stream operates at 2.5 Mbps (including overhead), the system can carry 15 digital video channels (for a total of 37.5 Mbps) in place of a single analog video channel. Since CATV systems typically provide in excess of 100 analog video channels, use of digital technology would permit 1500 digital video channels to be transmitted. If IPTV were being used for each of those channels, then 1500 simultaneous viewers could be supported. Since not all televisions are likely to be on simultaneously, this would be enough capacity to serve well in excess of 1000 homes, even if each home had more than one television.

Shared Data Networks

In many countries (for example, Japan, France, and Korea), high-speed data networks are already in place to serve individual subscribers. Also, many businesses have high-performance networks installed to many desktops. In this situation, it is a relatively straightforward process to overlay IPTV on an existing network that is shared with other data.

Adequate bandwidth is a key requirement for any type of video delivery system. For an IPTV system to work properly, each segment of the network must have enough capacity to handle the total bandwidth of all the simultaneous video streams passing over that segment. Multicasting (as discussed in Chapter 9) is one way to reduce the amount of bandwidth needed. In this case, the service provider's (or the corporation's) networking equipment fills the role of the DSLAM for performing the multicast "Join" and "Leave" commands required for channel changing. For VOD and VCR-like interactivity, the network must have enough bandwidth to handle the combined load of the individual streams.

Overall, it is quite feasible to implement IPTV services on a variety of different shared networks. A number of active implementations are in service today. The key to a successful IPTV deployment is ensuring that adequate capacity is available to service the peak number of simultaneous viewers, in terms of both video server capability and network capacity.

TECHNOLOGIES FOR IPTV

Various technologies are required to implement a practical IPTV system. Many are common to other video-over-IP technologies discussed elsewhere in this book; however, several are used primarily for IPTV applications. In the following sections we'll begin with a description of the systems used to prepare incoming content for IPTV distribution. We'll follow this with a deeper look at video servers. We'll conclude with a look at Multimedia Home Platform (MHP) and OpenCable Application Platform (OCAP), which are key technologies that support subscriber interactivity and the OSS functions required by service providers.

Content Processing

Content-processing systems accept real-time video signals from a wide variety of sources and shape them into a consistent format so that the customer STBs can decode the signals and display them on the viewers' televisions. This process involves several functions, as described in Table 13-3.

TABLE 13-3

Content Processor Functions

-
- *Compression*: For analog video sources, digital compression is performed on each video signal before it can be transmitted on an IPTV system. Items such as peak video data rate and packet length are made consistent between all of the different video sources to simplify the job of the transport and multiplexing functions.
 - *Transcoding*: For video sources already in a digital format, it is sometimes necessary to convert the MPEG profile or level of the incoming stream to one that is compatible with the STBs. For HD content, it is standard procedure to change MPEG-2 sources into H.264 to achieve lower bandwidths for DSL networks.
 - *Transrating*: Essentially, transrating is the process of changing the bit rate of an incoming digital video stream. For example, an SD feed at 4.5 Mbps may need to be reduced to 2.5 Mbps for use in the IPTV system.
 - *Program identification*: Each video stream needs to be uniquely labeled within the IPTV system so that multiplexing devices and user STBs can locate the correct streams. Each audio or video program within each MPEG transport stream (using the PID mechanism discussed in Chapter 6) must be processed to ensure that there are no conflicting PID numbers. If there are duplicates, then one of the streams must be renumbered to a different PID, and the associated PMTs must be updated.
-

Content processing can be performed on either a live stream or content that is to be stored inside a video server for later playback. When content processing is done on a live stream already in a digital format, the process is called *digital turnaround*. When this process is performed on content that is to be stored on a server before it is delivered to viewers, it is called *ingest*. This process also needs to capture the rules for using the content, such as limits to viewing duration and permissions for downloading. Real-time ingest is used, for example, to take a copy of a live video signal and store it for later viewing by subscribers. Off-line ingest is used for content that is delivered in tape, disk, or file form, such as Hollywood movies.

Video-on-Demand and Video Servers

Let's look at the architecture of a typical VOD system that uses video-over-IP technology, as shown in Figure 13-5. There are four main components to a VOD system. First, the content must be prepared for storage and delivery by compressing it and (usually) encrypting it on a content-preparation station. A VOD server




FIGURE 13-5 Typical VOD System Architecture

stores the content and creates the streams that are sent to viewers. Each viewer has an STB (or properly equipped PC) that receives the content, decrypts it, and generates a signal to feed the viewer's display. The STB also provides a means for viewers to order VOD services from the fourth and final element of the system shown: the subscriber management and conditional access system. This subsystem takes commands from viewers, sends appropriate commands to the VOD server, and delivers decryption keys to the STBs.

Video servers are essential to any VOD system, because they create the actual video streams that are sent out to each viewer. These servers can range in size from very large to fairly small, and they can be used for a variety of applications. In this section, we'll look at some of the different aspects of servers and how they are used for delivering content.


The amount of storage in a server can be large or small, and the number of streams supported can be large or small. These are not correlated; it is perfectly sensible to have a server with lots of storage and little streaming capacity if it is being used to hold video content that is only rarely viewed. Conversely, it is also sensible to have a

server with relatively little storage (say, 50–100 hours of video content) but very high stream capacity if it is being used to serve first-run Hollywood movies to many viewers simultaneously.

Many varieties of video servers are available. When purchasing a server, it is important to match the capabilities of the server to the task that needs to be performed. Video servers can be broken roughly into three categories:

- Production servers are used in the video production business, such as television networks and postproduction houses. For these customers, a video server must handle a great deal of content, in a variety of formats, and rapidly deliver files containing the content to user workstations when it is needed. These servers typically do very little, if any, streaming. Instead, the focus on these devices is large capacity and good support for content searching, including tools that support the use of meta-data and that can track multiple versions of files as content moves through the production process.
- Personal and corporate servers are used in environments where a relatively few streams need to be delivered simultaneously and the total amount of content is relatively low, such as a training department of a medium-sized corporation with a few dozen titles and less than 5–10 simultaneous viewers. This class of server can often be built with off-the-shelf components on a PC chassis with specialized software.
- Service providers need specially designed servers capable of storing possibly thousands of hours of content and delivering video streams to hundreds or thousands of simultaneous viewers. The capacity of these systems is truly staggering; in order to supply 1000 simultaneous users each with a 2.5-Mbps stream, the server needs to be able to pump out 2.5 gigabits of data every second. These units typically spread content across a large number of disk drives and use multiple processors in parallel with high-speed backplanes to format streams and deliver the content.

One important constraint on a streaming video server is that end-user devices (STBs) have a precisely defined amount of buffer available. If too much data is sent (causing a buffer overflow) or too little

**FIGURE 13-6** Centralized vs. Decentralized Video Servers

(causing underflow), the video signal to the viewer will be interrupted. To prevent this, the video server must be carefully designed to create streams that are well behaved and won't violate the buffer constraints.

Service providers use two main philosophies for server distribution in their networks, as shown in Figure 13-6. The first is centralized, where large, high-capacity servers are installed in central locations and the streams for each viewer are delivered over high-speed communications links to each local service provider facility. The second is decentralized, where smaller servers are located at each facility and provide streams only to local viewers. A central library server downloads copies of content to the distributed servers whenever necessary.

On one hand, the decentralized concept makes sense because it helps to reduce the amount of bandwidth needed between locations. On the other hand, the centralized concept makes sense because it reduces the number of servers that must be installed and reduces the costs of transporting and storing copies of the content in each of the different locations. Overall, both centralized and decentralized systems are used in practice, depending on system architecture, capabilities, and user viewing habits that affect VOD traffic patterns.

MHP and OCAP

MHP, which stands for Multimedia Home Platform, and OCAP, which stands for OpenCable Application Platform, are both software interface standards that have been developed for STBs. The Digital Video Broadcasting (DVB) Project, based in Europe, developed MHP; and CableLabs, a joint research and development group formed by members of the CATV industry in the United States, developed OCAP. To a great extent, OCAP is based on MHP. Both are open standards, although some license fees may apply.

MHP functions as a standardized interface for “middleware.” It provides a common interface between software applications and lower-level software, such as operating systems and device drivers. It does not provide user interfaces, but it does provide a platform on which user interfaces can be built. It is not an operating system; it defines a standard way that applications can access services provided by an operating system.

MHP greatly simplifies the task of application designers, by giving them a uniform platform to build on. Once an application has been designed and tested to run on MHP, it should be able to run on MHP implemented on any STB from different manufacturers. This capability helps provide a greater market for these applications than what would be available if the application was tied to the products of a single STB vendor.

MHP also simplifies the tasks of STB vendors, because it allows them to deploy applications written by third parties. This can help reduce the software development costs for a new STB and allow products to reach the market more quickly.

Table 13-4 defines some of the applications that MHP was designed to support. Table 13-5 defines some of the system resources inside an MHP device that applications can use to provide their intended functions.

Overall, MHP is designed to be a common interface for user interactivity applications and STBs, and gains in economy and efficiency should be possible if this standard is widely adopted. Perhaps the only drawback to this standard is the need for a fairly

TABLE 13-4

MHP Application Examples

-
- User interface for electronic program guide/interactive program guide
 - Home shopping
 - Access to broadcast information services—news, weather, financial information (super-teletext)
 - Video games
 - Applications synchronized to video content, such as sports scores, play-along games, audience voting/feedback
 - Secure transactions—home banking
-

TABLE 13-5

Examples of Resources Available to MHP Applications

-
- MPEG stream control
 - Text and graphics screen overlays
 - Video and audio clip playback
 - Program selector/tuner
 - Communication protocols (e.g., TCP, UDP)
 - Conditional access and encryption
 - User mouse/keyboard/infrared remote control access
 - Java virtual machine scripting
 - Media storage and playback controls
 - Internet access
-

powerful processor and a fair amount of memory to implement these functions in a low-cost STB. In the long run, anything like MHP or OCAP that can help drive software costs lower and provide truly open interfaces is good for service providers and, ultimately, consumers.

IPTV CASE STUDY

Kaplan Telephone is a local service provider located in south central Louisiana that has a long tradition of providing standard telephone service to a local area. Kaplan selected IPTV-over-DSL technology to offer video services to its customers, since it eliminates the need to build a second network alongside the existing twisted-pair

network. DSL also helped reduce the up-front capital expense, because subscriber equipment in the home and the central office was installed only when subscribers actually purchased service. Also, because DSL can provide data services alongside traditional subscriber line voice services, Kaplan can provide a full "triple play" of consumer services.

For the digital video headend in 2004, Kaplan selected Tut Systems of Lake Oswego, Oregon,¹ who supplied the Astria[®] Content Processor. To understand the required functions, let's look at how the different incoming signals are processed.

- Satellite signals normally arrive in digital format, often scrambled or encrypted to prevent unauthorized reception.² The satellite receivers are responsible for processing the signals from the satellite dish; they can either provide streams in DVB-ASI format or convert one or more programs into baseband digital or analog video signals. Receiver processing includes any necessary descrambling or decrypting of the video signals.
- Local off-air video signals can arrive as either digital or analog signals. The demodulator is responsible for taking the broadcast video signal and converting it into either a baseband analog or digital video signal.
- In many cases, video service providers such as Kaplan are permitted to insert commercials into certain television channels at certain times. These local ad spots can come from a variety of sources: videotapes, DVDs, file transfers, etc. The spots are stored on a server, ready to play whenever they are required. The output of the server is fed into a monitoring switcher, where the commercials can be spliced into the video program streams as needed. Increasingly, digital ad splicing is being utilized that keeps both the primary stream and the ad content in the compressed domain.
- A legal requirement in the United States is the broadcast of Emergency Alert System (EAS) signals; similar requirements are present in many other countries. The EAS provides information generated by official government agencies to television

1. Tut Systems has since been acquired by Motorola.

2. Analog signals are still sometimes used on satellites, but they are becoming increasingly rare.

viewers regarding severe weather events and other emergency situations. Local television providers are required to broadcast the EAS signals on the television channels they are distributing. Many times, the EAS messages are text based and inserted into the bottom portion of the video screen. Suitable messages are created inside the character generator and fed into the MPEG encoders for the broadcast video channels.

Figure 13-7 shows a number of different functional blocks inside the Astria® Content Processor-based headend. On the input side, demodulators and receivers are used to convert incoming signals from their native format (broadcast television or satellite signals

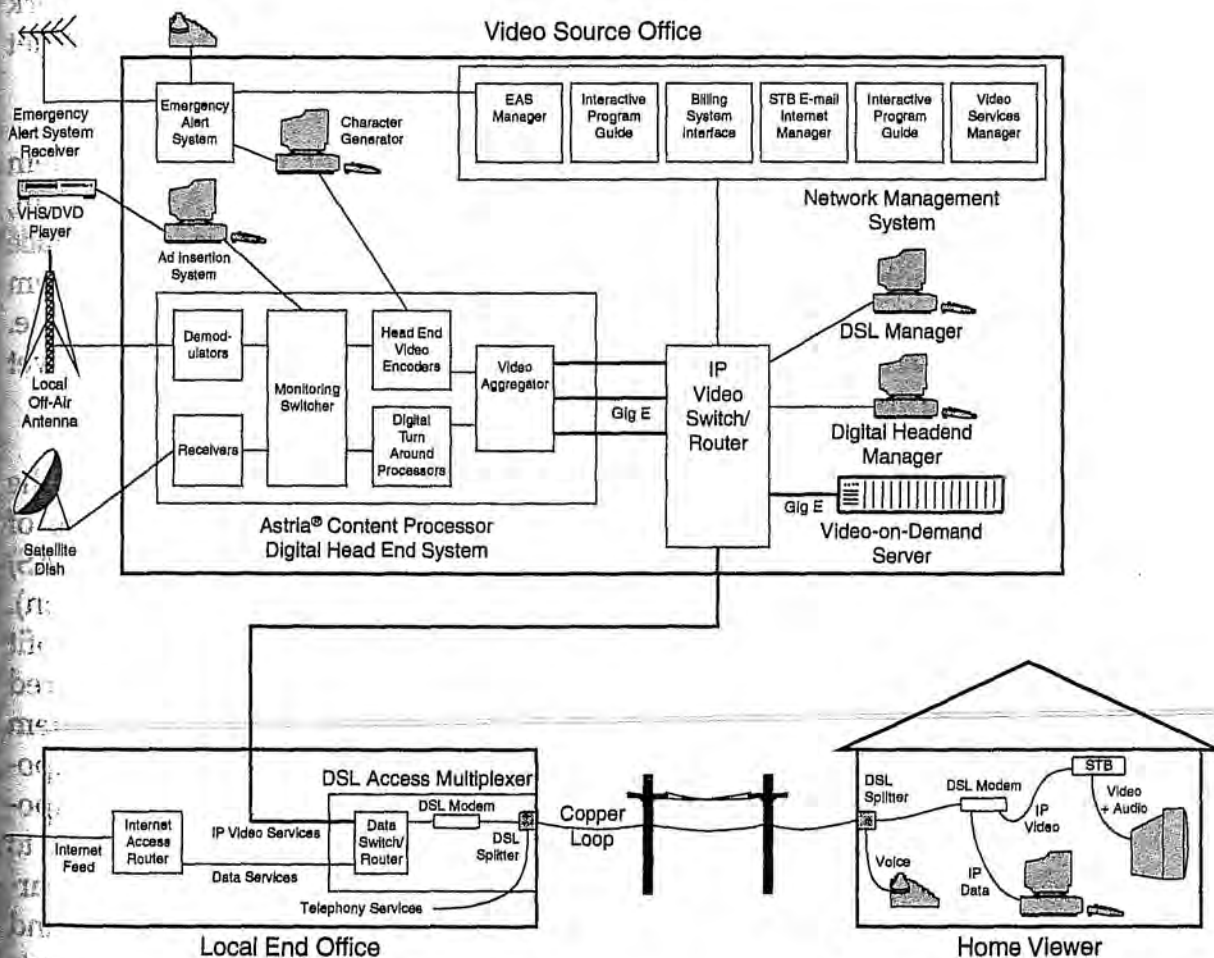


FIGURE 13-7 DSL Video System Example

with various modulation schemes) into video signals that can be processed. These signals are fed into a monitoring switcher that allows local advertisements to be switched or digitally spliced into selected video streams and provides a method for personnel to tap into video streams as needed to observe their performance. Some of the signals flowing out of the monitoring switcher are baseband video; these signals need to be digitized and compressed by the headend video encoders. Other signals that flow out of the monitoring switcher are in digital format, so these are processed to make them compatible with the DSL distribution network, by means of transrating or transcoding performed by digital turn-around processors. The final function is called *video aggregation*, in which many different video streams are multiplexed into large-capacity (gigabit Ethernet) outputs. These signals flow out to the various DSL access multiplexers around the Kaplan network by way of high-speed switches/routers and a gigabit Ethernet distribution network.

The main Operations Support System (OSS) for the Kaplan system covers a variety of functions, including the Emergency Alert System, the various video services offered to subscribers, the data interface to the subscriber billing system, the e-mail/Internet access system for subscriber set-top boxes, and the interactive program guide. A separate system is used to manage the DSL hardware and yet another to manage the Astria[®] Content Processor.

One point to keep in mind when reviewing the Kaplan system is the variety of different functions that must be accomplished, some of which are mandated by government regulations (such as the EAS) and others driven by economic realities (such as local ad insertion). For an IPTV system to work properly and create an operating profit for the system owner, these various functions need to be delivered efficiently and cost effectively by the system provider(s). System owners have to consider not only the costs of the individual components, but also the costs associated with integrating all the components to form a working combined network. This is in addition to the costs of purchasing the rights to content and the costs of preparing the content for distribution. Typically, the costs of acquiring and installing all of the necessary hardware and software components can be substantial; these costs need to be amortized over a number

of years with a substantial number of paying subscribers before economic breakeven can be reached. Modeling these cash flows can be tricky; however, it is an essential part of IPTV system planning.

REVIEW AND CHECKLIST UPDATE

In this chapter, we covered IPTV services to the home. We began by looking at some of the market forces that are making a business out of this technology. We discussed video-on-demand (VOD) and some of its various flavors. Then we examined the functions of the Operations Support Systems (OSS) that are crucial to large-scale deployment of any kind of video to the home system. A few applications were described, and we explored some of the key technologies for IPTV, including DSL, PON, and CATV transport; content processors; video servers; and the innovative middleware technology known as MHP and OCAP. Finally, we investigated a typical IPTV installation.

Chapter 13 Checklist Update

- ☐ Develop a business plan prior to deploying IPTV to the home. Make sure to include capital expenditures (hardware and software), installation costs, programming costs, and ongoing maintenance and customer support costs.
- ☐ Determine the types of viewer interactivity that will be supported.
- ☐ Determine how live content will be delivered to viewers. If digital turnaround will be used, determine the format of the incoming digital signals.
- ☐ Consider how this might change in coming years as HD content becomes more widespread and as content owners change their requirements.
- ☐ If video-on-demand is to be used, consider the following issues:
 - ☐ Select a method for delivering VOD services: NVOD, pay-per-view VOD, or subscription VOD. Make sure that systems are in place to support these functions if they are basic, VCR-like, or content interactive.
 - ☐ For large installations, select a video server deployment strategy: centralized, distributed, or a blend of the two.

- ☐ Calculate the total number of hours of programming that will need to be stored on each server, and ensure that adequate storage capacity is purchased.
- ☐ Determine how many simultaneous streams each server must deliver, and ensure that each server purchased has enough stream processing capacity.
- ☐ Make sure an ingest system is available to get prerecorded content from a variety of sources onto the video servers.
- ☐ Select a supplier of electronic program guide data.
- ☐ For user software platforms (STBs and the systems that support them), consider using open-standard middleware, such as MHP or OCAP.

15

INTERNET VIDEO

Internet video is a relatively new phenomenon in the world of video delivery. Just a few years ago, people who wanted to send video over an IP network had to build their own infrastructure, including video encoders, delivery systems, client playback hardware/software, and a management system to connect viewers with the content they desire. Plus, due to performance constraints, most delivery had to be done over private networks. Now, all or most of that groundwork has been done by a variety of online services and off-line utilities, so the costs of entering this market have dropped dramatically. In addition, millions of viewers have upgraded to broadband Internet connections and have become accustomed to searching for videos and playing them through web-browser plug-ins already downloaded and installed on their client devices.

YouTube is by far the most viewed site on the Internet for video content. While there is a significant amount of junk, there is also a surprising amount of high-quality content produced by talented individuals. In addition, there is an increasing amount of professionally

produced content supported by paid sponsors. Some of these latter videos are hosted on YouTube but embedded in corporate web pages for viewing.

Many users will have encountered video streaming on sites other than YouTube while surfing the web, particularly if they sought out video content. Popular video sites include CNN.com for news stories; sonypictures.com, warnerbros.com, and Disney.com for trailers of upcoming movies; and mtv.com and music.yahoo.com for music videos. In addition, a number of sites offer continuous channels of television-quality programming over the Internet, including researchchannel.com (a lovely 1200-kbps signal) and Bloomberg television (at www.bloomberg.com).

In this chapter, we begin with a discussion of the key attributes of Internet video, including the basic system architecture. We'll then go into some of the unique applications of Internet video. This is followed by an overview of a typical system architecture and different platforms that can be used for Internet video services. We'll wrap up the chapter with a discussion about some specific technologies, including client software and embedding videos that are hosted by other websites.

KEY ATTRIBUTES OF INTERNET VIDEO

Viewers' experiences and expectations for Internet video are very different from those for most other forms of video entertainment, such as IPTV and broadcast television. Most viewers have much lower expectations for Internet video, particularly if they have ever tried to watch video over a dial-up connection. People seem willing to tolerate things they would never tolerate with broadcast television, such as seconds (or even tens of seconds) of blank screens before a video starts to play, low resolutions on many websites, issues with video and audio synchronization, and other technical drawbacks. Of course, technology marches on, and the Internet video viewer experience continues to improve.

Figure 15-1 shows a very simplified view of an Internet video network. This diagram is broken into two sections, labeled *production*

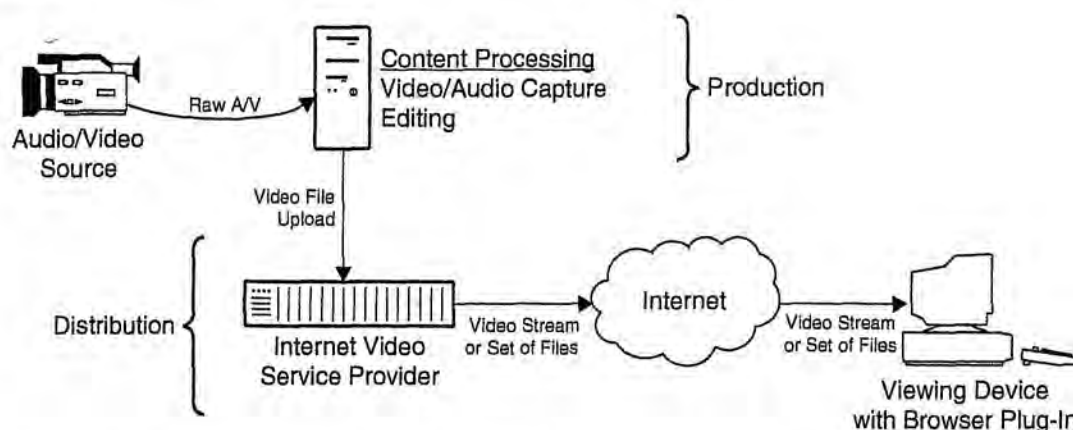


FIGURE 15-1 Typical Internet Video Network

and *distribution*. In production, the video content is captured from a source, digitized, edited, labeled, and bundled into a file that is placed on a server where it can be accessed. In distribution, a person uses an Internet-connected PC to search for content, connect to the server, acquire rights to view the content, and then either download a video file or request a video stream of the content for viewing on their PC using specialized multimedia viewing software.

A person using a PC or other device initiates each viewing session. First, the user must identify where the content is located on the Internet. For example, a user might have received an e-mail from a friend with a link to a website containing the video. When the user clicks on the link, the browser on the PC connects to the appropriate web server. Typically, the web server then displays a screen that gives some information about the video (such as a description and the clip's duration). The user then may be asked to click on a link embedded in that page, which begins the video playing process. One important step that happens at this time is an exchange of information between the server and the browser software on the PC that tells the browser to run a plug-in application called a *media player*. The media player is fed data that the browser has retrieved from the web, decodes the video, and converts it into an image that can be displayed. If the proper media player software isn't installed on the PC, the user is normally prompted to install it by downloading the software from a suitable source. Then, as the video file is delivered, the user can watch the content.

We first defined the key attributes of Internet video in Chapter 1: discrete content elements that are selected by users from millions of choices that are delivered via the Internet in a wide variety of formats for playback using PCs or other devices. This is in contrast to IPTV, which provides hundreds of continuous streams of professionally produced TV channels that are delivered over dedicated networks to televisions by means of a set-top box. Internet video should also not be confused with podcasting, where video content files are delivered to portable devices for later playback and user storage. Both IPTV and podcasting are covered in other chapters in this book. Let's discuss the attributes of Internet video in more detail.

Discrete Content Elements

Instead of continuous channels of highly produced programming, most Internet video content is available in the form of clips for viewers to watch at a convenient time. Many of the video files available for viewing or downloading are relatively short, five minutes or less. (YouTube imposes a limit of 10 minutes on most of the video clips that it hosts.) Certainly, longer-duration files are available, but they tend to be available on advertiser-supported or viewer-paid sites.

Of course, not all Internet video involves discrete content elements; some real-time streaming broadcasts are available. For example, NASA TV offers some live video content each day from the International Space Station as well as live coverage of major events such as shuttle launches and space walks. In between, educational, news, and other programming are provided. For-profit real-time Internet video channels are also becoming a reality, as the number of viewers with high-bandwidth Internet connections reaches a level that is attractive for subscription or advertising-based services.

Millions of Content Offerings

Any quick search of some of the more popular video websites will show that there are literally millions of different video files available for viewing, with thousands more being added each day. Locating a specific piece of content for viewing can be a challenge with Internet

video. Many viewers find content by following links on web pages or in e-mails that direct them to video content sites. Others use the listings of popular titles on these websites. Still other viewers find videos to watch using general-purpose search engines (such as Google) or site-specific search engines. Unfortunately, there is no master program guide for the Internet—there's simply too much new content being added each day for this to be practical.

Multiple Content Formats

A wide variety of formats can be used for video files, and virtually all of them have found their way onto the Internet. There are many choices, including various camera formats (such as DV), MPEG family (1, 2, or 4), JPEG (basic or 2000), player-specific formats (Windows Media, QuickTime, Real Networks, etc.), and a variety of computer file formats (such as AVI). Consumers who view a significant amount of Internet video content often end up with a variety of different video players loaded onto their computers in order to handle the various video file formats.

For content providers, this variety can present a dilemma. If the provider chooses to support only a single video format, then any viewer who wishes to watch the content must already have the appropriate player software installed or find a way to get the proper player (most of which are distributed for free). If, on the other hand, the provider chooses to support multiple formats, then they assume the burden of producing and managing multiple versions of their content.

In addition to the choice of video compression technology, content providers must decide on the screen resolutions they will support. Low resolutions offer small file sizes that are easier to download over low-bandwidth network connections but also create small images for viewing. Higher resolutions offer increased picture quality but can require a long time to download or a high-bandwidth connection for live streaming. Of course, the choices may be limited, since some devices (such as Apple's iPod) and some video-hosting websites (such as YouTube) only support specific image resolutions.

Delivered Over the Public Internet

One big strength of Internet video is that it can be delivered to any viewer with a connection to the Internet. Of course, high-bandwidth connections are easier to use and deliver quicker results, but even consumers with low-speed dial-up connections can download video files if they are patient enough.

Because video sites can be accessed from around the globe, the potential audience for any video can be very large, particularly if a good mechanism is in place to inform consumers about available content.

Use of the public Internet also means that content providers don't need to build network facilities to their viewers, resulting in a significant cost savings. Unfortunately, this means that the network must be shared with a host of other applications that consume bandwidth. Also, there is no means for video content to be given higher priority than other types of traffic, which can dramatically increase the difficulty of delivering high-quality, high-bandwidth content in real time to viewers, as is commonly done on IPTV systems.

Viewed on Consumer PCs

A reasonably powerful consumer desktop or laptop PC is capable of running the player software required to decompress and display most compressed video formats. Performance can sometimes be helped through the use of graphics accelerator cards or additional memory that is added to the system. In some cases, viewers will watch the content on the display screen of the PC itself; in other cases, the video will be displayed on a television set connected to a video output port of the PC.

Other consumer video playback devices have begun to enter the market for Internet video content. One of the most popular portable video viewers is the Apple Video iPod, which features a screen resolution of 320 × 240 pixels. Most of these portable devices have a limited range of video file types they will support, so it is essential for consumers to select only those content files that are compatible

with their device's capabilities. Some appliances have also appeared on the market that can receive video content directly from an Internet connection and display it on a television.

INTERNET VIDEO APPLICATIONS

Internet video can be used for essentially all the same applications as other video-delivery systems, including entertainment, advertising, videoconferencing, and training and education. However, there are some unique applications for Internet video that bear discussing.

Narrowcasting

In narrowcasting, IP video signals are sent over an established data network to reach an audience for programming that might not appeal to a large audience. For example, for the past five years the PGA Tour has broadcast live Internet video coverage of every player on a single hole during the first two days of a golf tournament (the island 17th green of the Tournament Players Club at Sawgrass course used for The Player's Championship tournament, for those who must know). Also available are major league baseball games, television feeds from a number of different countries, and a great deal of prerecorded content. A surprising amount of content is available by narrowcast; consumers need only go to the homepages of any of the major media player companies to see a smattering of what is available.

Narrowcast video streams are normally displayed only on a user's PC, but the same signal can be processed for display on a normal television as well, using a PC equipped with a video output card or a stand-alone adapter. Narrowcasting content is provided free of charge (supported by advertising), on a pay-per-view basis, or on a subscription basis (such as RealNetworks' SuperPass).

A typical narrowcasting system is set up using a centralized server for video content. Viewers obtain the video by connecting

to the server through the Internet or other IP communications network. The streams can be served using a streaming protocol such as RTP or a download-and-play protocol such as HTTP. Because the Internet is not multicast enabled, the server needs to create a stream for each user. For unexpectedly popular live services, viewers can experience choppy or even frozen video if the servers don't have enough capacity to meet the surge in demand. To prevent these problems, an admission control method is often used to deny service to newcomers once system capacity has been reached.

Movie Previews

Video over the Internet really got a jump start in the late 1990s with the availability of Hollywood movie previews on websites. No longer constrained to showing their wares on expensive television commercials or as "coming attractions" in theaters, movie studios began to reach out to the online community. At first, much of this content was set up for download and play, because very few home users had access to broadband connections suitable for high-quality streaming. One of the classics was the *Star Wars Episode 1* trailer at a (then) hefty 10.4 megabytes; it was released in March 1999 and downloaded 3.5 million times in its first five days of availability, according to a press release at the time.¹ Many of the downloads would have taken upwards of half an hour for people who, like the author at the time, had dial-up service. (A good source of movie information and links to trailers can be found at www.imdb.com, the Internet Movie Database.)

By the end of 2007, over 80 percent of the active Internet users in the United States employed a broadband service (cable modem, DSL, wireless, or similar). As a result, more and more content is available in streaming form. Today, a user can log onto a number of different websites and look at movie trailers from essentially all new and forthcoming Hollywood releases.

1. From a press release that was previously available in April 2004 at <http://www.starwars.com/episode-i/news/1999/03/news19990316b.html>.

Internet Radio and TV

A number of free and subscription services have appeared on the Internet to provide both audio and video content. There are literally thousands of Internet radio stations, due in part to the relatively low cost of equipment and the low bandwidth required. Internet television stations are much less common, but they are becoming more feasible as the number of users with broadband connections increases.

Video- and audio-streaming sites have been developed for a number of purposes, including corporate branding (free sites), advertising supported (also free), and subscription (monthly or other periodic payment system). A huge variety of content is available, including newscasts, music, adult programming, and entertainment. Because this material is organized similar to a traditional radio or television broadcast, users are restricted to viewing the content in the order in which it is presented. It is not a content-on-demand service, where each user can watch any content he or she chooses in any order.

Much of this content is prerecorded, so download-and-play technology is perfectly adequate. However, this can be somewhat disruptive to the flow of the broadcast, because each file must be downloaded before play can begin. (This technology is much more disruptive for video files than for audio files, simply because video files are much larger and take much longer to download.) Progressive download greatly alleviates this problem, because playback of each new file of content can begin as soon as the first segment of it is downloaded to the PC.

Live Content

For live content, particularly for any type of interactive Internet video, true streaming is the only practical choice. When this is implemented, the video signal goes directly from the source into a real-time video encoder, and then output is immediately wrapped in IP packets and transmitted out to the network. Adequate bandwidth needs to be available along the entire path from the source to the viewer for the entire stream, including video, audio, and any control

data. At the viewer's location, the client device receives the incoming stream, decodes it in real time (often by use of a browser plug-in), and delivers the signal directly to the viewing device.

Webcams, which are now embedded in many laptops, are popular sources for live content, particularly in two-way connections. Similar in functionality to video conferences, these applications use different technologies to achieve the same goal: live two-way communication with combined video and audio signals. As long as the delay is short enough, almost any Internet video architecture can be used to support webcams, including P2P networks such as Skype, which has built-in support for several types of video camera. In many applications, the bandwidth of webcam video is kept low by using low frame rates and low video resolutions. This also helps limit the amount of processing power needed for the compression and decompression calculations and reduces the amount of internal device bandwidth occupied by the raw and compressed video signals. These considerations are not trivial, because any PC involved in a two-way webcam call must simultaneously encode an outgoing audio/video signal and decode an incoming audio/video signal, which can create a significant processing load.

Another popular application for live Internet video is live viewing of traffic and weather conditions at various locations around a metropolis. These are typically configured to support a number of simultaneous viewers, possibly through the use of a content-delivery network (CDN, described later in this chapter) if large numbers of viewers are to be supported simultaneously. Enough processing power is needed to generate a unique IP stream for each viewer, with some mechanism required to handle overload conditions, which may occur, for example, during periods of severe weather.

Placeshifting

Many readers may be familiar with the concept of timeshifting, where viewers record off-air video programs for playback at a more convenient time, usually by means of a PVR. *Placeshifting* is a related concept where the viewer is in a location outside the home but still able to watch video programming available at home. This

is accomplished through the use of a small Internet streaming server located at the home that is connected to the viewer's normal video equipment. The user can remotely control this server by means of their PC that is also used to display the content. In addition, the user can control the home video equipment by way of software controls on the PC that are converted into infrared remote commands in the home installation. Sling Media is a pioneer in this market space.

One of the big benefits of this technology is enabling users to watch subscription TV programming that they have available in their home but not in their current location. Local news and sports programming might also be of interest. In addition, with the right connections and software, the remote user could view programming stored on their home PVR or other video content storage device. Of course, there are downsides. This technology requires use of the valuable return path bandwidth from the home to the viewer's ISP, which typically has limited bandwidth. In addition, any devices such as CATV or DTH STBs that are being used and controlled by the remote viewer are not available for viewers in the home to use, unless they are watching the same programming. Overall, placeshifting is an interesting concept, which may become popular with some classes of viewers.

SYSTEM ARCHITECTURE

The basic architecture for an Internet video system is essentially identical to the ones used for private video streaming or any web-based service. A server loaded with content is able to deliver files to client devices upon request. These client devices take the incoming media content files and decode them using a software application for presentation to the user via a connected display.

There are some differences, of course. An Internet video server has a fixed public IP address that is made visible to the World Wide Web and is accessible to any connected user. Content-delivery networks (CDNs) can be used to assist in the delivery of files, particularly for popular websites with large amounts of traffic that could easily overwhelm even a high-capacity web server. Instead of a program guide that many private video systems provide, users can browse for content by going to a suitable search engine or simply click on links embedded

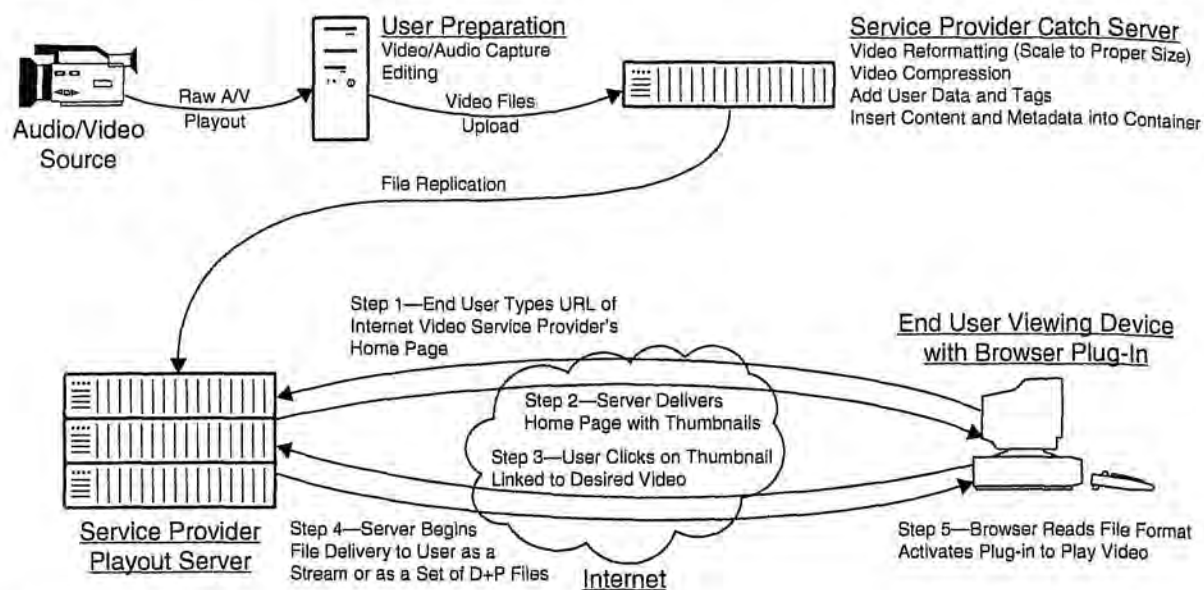


FIGURE 15-2 Detailed View of Internet Video Network

in information sources they trust. Web browsers with video decoding plug-ins are used in place of stand-alone viewing software. Figure 15-2 shows a more complex view of an Internet video system.

Many video delivery services today use progressive download and play for one big reason: the ability to transit firewalls. Many firewalls block all or most RTP/UDP packets to prevent rogue users from inserting packets containing viruses or other malicious code into legitimate content streams. (TCP streams are much more immune to these types of attacks because of the packet sequence numbers used to ensure that all packets are delivered properly.) The files used in progressive D+P applications are essentially identical to any other file that can be delivered by HTTP, and in fact these files can be served by any standard web server software (such as Apache).

The key to smooth playback of the video signal is to ensure that each video file is delivered before it is needed. To make this easier, many players establish a buffer with enough storage to hold, say, 15 or 30 seconds' worth of video before playback begins. Then, during playback, the player continues to download more content files in sequence, normally trying to maintain a constant buffer of content that is ready to play. In the event that delays occur in the download,

the buffer may shrink; but as long as it doesn't become empty, video playback will continue smoothly. If the buffer does become empty, then the video display will typically freeze until more content can be downloaded and the buffer becomes at least partially filled. This happens noticeably more often on HD content than on SD content because of the larger amounts of data that must be streamed for HD.

For those delivery services that do not use progressive D+P, video streaming is another popular choice for distributing video content. This technology is required for live events and can also be used for applications that need extremely fast start-up or large amounts of user playback control. One big advantage of true streaming is that a video file is never stored on the user's device—it simply occupies some memory for buffering during payout. Of course, if the connection between the server and the viewer is interrupted, even momentarily, the video signal will be corrupted, unless a large enough buffer is used.

DELIVERY PLATFORMS AND SERVICES

A number of delivery platforms are available for use in Internet video applications. They can be broken down into three rough categories: free services, free services with in-stream advertisements, and user-paid services. *Free* services are ones that a viewer does not need to pay for but may include advertising that is not part of the actual video stream. *In-stream advertisement* services are ones that use video advertising inserted into the content, either before, during, or after the video the user actually wants to watch. *Paid* services are ones that the user pays for, either on a per-item basis or on a subscription basis. Of course, since there are many areas of overlap, some additional explanation of the types of advertising will be useful.

Types of Advertising

Many different ways can be used to communicate an advertiser's message to a viewer. These range from fairly noninvasive to quite invasive and can be combined on the same video at different points in the content file. Figure 15-3 shows the types of ads as they would appear in a media player display.



FIGURE 15-3 Types of Video Player Advertising

- *Website or banner ads*, which may be text or graphics that appear adjacent to the video display window. These may be static or animated and may or may not bear any relation to the content being played. Typically these ads include a hotlink that will cause another browser window to open up if the viewer clicks on them.
- *Animated overlays* are graphics that appear in the video window while the user's selected video is playing, but they do not interrupt the flow of the main window or the audio track. In many implementations, overlays appear a short while after the video has started playing, last 10 seconds or so, and then disappear if the user does not click on them. (Sometimes the video display window will show a button or a tab to reactivate the overlay after it disappears.) In most implementations, the overlays occupy only a portion of the video window (perhaps the lower 20 percent) and may or may not be translucent to show the video behind. If the user does click on the overlay, then either another browser window can be opened or the current video can pause to allow the advertiser's content to be displayed.
- *In-stream advertising* consists of video content that completely replaces the viewer's selected content for a short duration. This user experience is exactly the same as with traditional television advertising, where the flow of the program is interrupted and replaced with completely different audio and video. When the

ad has completed playing, the user's content will resume automatically. These ads can be placed before the viewer's content (called *preroll*), in the middle of the content (*mid-roll*), or after the viewer's content has played (*postroll*). These three different delivery scenarios bring obvious differences in viewer experience, so content suppliers need to choose their model carefully.

Free Services

Internet video delivery services are considered free when there is no direct payment by the user to access the video and they are also free of embedded advertising. This does not mean there is no advertising on the site; it simply means that the advertising is either in the form of website or banner ads. These ads are often placed selectively; advertisers generally do not want their ads to appear on a random collection of videos that may contain content of questionable origin or quality.

YouTube is the prime example of a free service. It provides an enormous library of user-generated content, some of which is in violation of content owners' rights. In addition, a significant amount of content is provided by professional sources, such as movie studios, music labels, and other companies trying to target a specific demographic for their products. There is even a surprising number of actual television commercials on the site, although these tend more toward humor. Profit for providers can come from programs such as YouTube's partner program, where the revenues generated by advertising that appears alongside videos during playback are shared with the providers.

Advertising-Supported Websites

Advertising-supported websites typically make use of in-stream advertising, often in the preroll format (advertising plays before the requested content plays). These sites offer primarily licensed content that is professionally produced for entertainment purposes, such as episodes or excerpts of television shows, movies, animated shorts, and any other content of known quality to both the website provider

and viewers. This level of trust is essential, because viewers would quickly tire of a website that did not provide good quality content after they have been forced to view a 15- to 30-second commercial.

Hulu is a good example of a website that provides a wide range of licensed video content coupled with mid-roll, in-stream advertising. In the playback timeline on the video player each of the ads is indicated by a small dot, enabling the viewer to see when the next ad will occur. A very nice feature of Hulu requires a viewer to watch a commercial only once; if the viewer rewinds after watching an ad, then the ad will not play a second time to that viewer in that session.

Paid Services

Paid services are almost entirely licensed and provide the most valuable forms of content, including live sports and adult programming. Essentially all the content on these sites is licensed from a provider; even supposed “amateur” adult video performers get paid for their appearances on these sites.

One example of a paid video site is MLB.com, which is owned by major league baseball in the United States. During the baseball season, sports fans can subscribe to the site, allowing them access to live video coverage of all of the televised baseball games not in their home city (so fans will still have incentive to attend games in person).

Content-Delivery Networks

CDNs such as Akamai and Limelight provide a range of valuable services to Internet video content providers. One of the basic functions is to deliver large media files to a user who has requested a web page from a content provider. Another function is to host reflecting servers that can be used to simulate multicast streaming over the Internet.

Figure 15-4 shows a basic transaction for a simple web page request that is supported by a CDN. In this example, the user has requested a web page that contains large files (such as a graphic or a video) from a website hosted by Taggart Transcontinental. This website

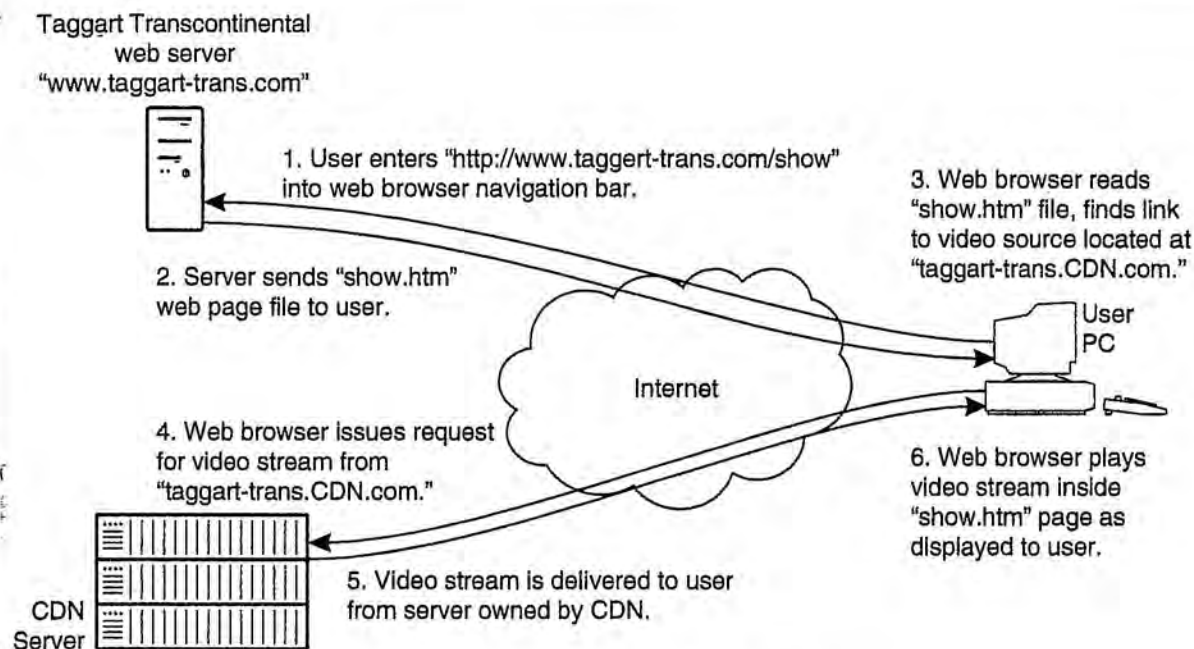


FIGURE 15-4 Typical Internet Content-Delivery Network Application

responds by sending a simple HTML file that contains the code for the overall shape and basic working of the web page, along with a series of links to the large media files. Since these files are not hosted by Taggart but, instead, by the CDN, the URLs in the HTML file point to addresses that belong to the CDN. The CDN then proceeds to deliver the files to the user's browser, where the content can be displayed. Live stream reflecting works in a similar manner: Browsers are redirected to obtain a live stream from a CDN server by means of a link embedded in an HTML file that is downloaded from the website advertising the stream.

CLIENT SOFTWARE

Since all Internet video is delivered in a compressed format, the viewer needs to have the correct software installed on an Internet-connected PC. Suitable software exists for a wide range of platforms and operating systems.

Viewers may or may not have their choice of media player software to use, particularly if DRM is employed. Most DRM systems are

closed; that is, content files that have been encoded with DRM will only play back on a specific brand of media player, and sometimes the player needs to be a specific version.

Browser Plug-ins

Some standard web browsers come equipped with little or no ability to display video files, and through the use of media player plug-ins, users can view a wide variety of content. Plug-ins (also known as *Browser Helper Objects* in Microsoft's Internet Explorer) contain software that is triggered by the browser when needed to display certain types of content, such as an Adobe Flash file. They can also be used to perform functions such as deciphering encrypted content and adding functions to a browser such as a toolbar.

One big advantage of plug-ins is that they can add significant functionality to a browser without requiring modifications to the browser as a whole. In particular, as video-streaming formats are updated or as encryption schemes are enhanced, changes can be implemented simply by updating a plug-in, not modifying the entire browser.

Many plug-ins have advanced capabilities beyond video decoding. They can present an entire customized user interface for playback, complete with advertisements, playback controls, etc. They can also restrict what a user can do, such as forcing a banner ad to be visible while the media playback is happening.

Stand-Alone Media Players

Many video decoders are available in the form of stand-alone media players, such as the QuickTime player. These stand-alone programs work like any other application on a PC; they do not depend on a web browser for operation. These players have a number of built-in capabilities, including the ability to play files stored on the local disk of the PC (often including DVD drives), the ability to locate and play files on the Internet, and the ability to play live streams from a number of sites. In addition, these players can launch other

programs, such as web browsers, to display other nonmedia content when the user navigates to normal web pages. The most common stand-alone players also offer users the ability to purchase and download content directly from the Internet, without having to go through a separate web browser interface.

From a content provider standpoint, the differences between plug-in and stand-alone media players are subtle. Plug-ins are more automatic for users; because plug-ins are loaded when the browser loads, they are ready to play as soon as the content arrives. Stand-alone players give viewers more control over how their media players appear, through the use of “skins” that change the borders around the viewing screen and the user controls. Stand-alone players also typically give the user more ability to control the size and location of the playout window, and they are easy to separate from a web page.

Content Preparation

Preparing content for Internet video distribution is a straightforward process. First, the video needs to be recorded and captured into a computer file. Editing can then take place, using any of a number of commercially available programs. Some basic editing packages are provided free with operating systems or cameras; others with advanced capabilities can be purchased from a number of suppliers. This is also the stage where sound, music, and effects should be added and any titles or credits should be prepared. This stage is done when the video is completely ready for viewing with all desired content included, because there is no practical way to add additional content in the subsequent steps.

The next step is to compress the video content. This step is required from a purely practical standpoint—uncompressed video files can be extremely large and would take a long time to upload to a website. In addition, most Internet video hosting sites have a limited number of video file formats they will support, and they may also have a limit on the size of the files that can be uploaded. For all these reasons, it makes sense to do a first-pass compression during the editing process.

Many Internet video-hosting sites will take uploaded videos and compress them into a standard format. This allows all of the (potentially millions of) content items to be stored and processed in a consistent manner on the site and for a uniform video player window to be used by all of the content. Depending on the site, more than one version of the compression can be done on uploaded content to provide low-bandwidth, low-resolution videos to users with low-speed Internet connections and high-resolution (even HD) videos to other users.

Some video producers have reported better video quality when video signals are compressed to exactly the target format before uploading. By selecting the correct file format, compression codec, resolution, aspect ratio, and frame rate, a user can match closely if not perfectly the desired format of the website.

A number of free and paid tools are available for performing editing and compression. As a rule, the free tools offer fewer editing features, generally saving more advanced scene transitions such as fades and wipes for paid versions. Also, the free tools offer a smaller selection of encoding formats and rates for both audio and video, and they don't offer the automation capabilities of the paid tools. However, the free tools do an adequate job for most content, and because they are free it is likely that a content developer will have access to a wide variety of them on an authoring system.

Inserting Video into Websites and Blogs

With the website development tools available today, it is possible to include many different types of content in web pages, such as photographs, illustrations, sound clips, animated buttons, and video content. From the viewpoint of a web browser, this multimedia content simply looks like any other content on the website—a series of files that need to be downloaded and displayed to the user. With the proper tools, HTML or XML code can be written to format the web pages, and any media files encountered by browsers are handed to plug-ins for user playback. So, in a basic sense, all that is needed to embed video in a website is a suitable website authoring tool.

When a single web server hosts all the different multimedia elements for a website, bandwidth charges will accrue whenever a visitor views the content. To help mitigate these charges, it is possible to upload content to a video-hosting website such as YouTube and still have it appear to be part of the website. When this process works, the website owner is responsible for the bandwidth charges of the materials hosted on the main server, but the video-hosting website is responsible for the bandwidth needed to deliver the hosted content to the viewer. This process, called *embedding*, is supported by some video-hosting websites. Note that embedding may also be preferred to hotlinking because it helps to prevent visitors from leaving the original website and never returning. There is typically no charge for this, but the hosting site may insert their corporate logo into the video player window. Table 15-1 lists some of the activities supported by video-hosting websites for use by independent website developers.

TABLE 15-1

Activities Supported by Video-Hosting Websites

-
- *Video uploading*: The process required to upload videos into a site like YouTube can be daunting for many web users. Using APIs (application programming interfaces) supplied by video-hosting websites, it becomes possible for an independent website developer to allow his or her visitors to upload videos to the hosting site without having to navigate to the hosting site's homepage, log in, and upload the video, greatly simplifying the process. This tool also allows the independent website to add tags to any videos that are uploaded to allow later searching based on these tags.
 - *Displaying a list of videos*: Some website developers might wish to provide visitors with a list of different content options to view without leaving the current page. This capability can allow a website to provide a sort of program guide to visitors.
 - *Controlling an embedded playback window*: Managing the look and feel of a website is important to many developers, so video-hosting websites have responded by allowing independent developers to design their own user interfaces for the playback window. Using APIs for JavaScript and Flash Player, the website developer can change how the video window borders appear on the screen, start videos at specific times, and synchronize other events on the webpage, to name a few tools.
 - *Add or remove playback controls*: The user controls can also be extensively modified, to provide, for example, an oversize audio mute control for videos that may have dynamic audio performances.
 - *Automatic blog posting*: For bloggers who wish to embed video content into their blogs, the process of doing so can be automated through the use of tools on some video-hosting websites. With these tools, the blogger can simply log in, select the video(s) to be posted, and then proceed to add any comments or other material he or she wants.
-

Finding Content

There are a number of ways to locate content on the Internet. Normal search engines (such as Ask.com) often lead to web pages that have video content, based either on the titles that users have given the videos or on tags that have been associated with the videos by their owner. Specialized search engines are also available, such as video.google.com, which will search only for video content on a wide variety of web pages, including many not hosted by Google or YouTube. Searching on websites specifically devoted to video content (such as YouTube) can be done either with keywords, as with a normal web search, or through the use of tags. Users who upload videos can apply tags in the form of keywords to their content to make it easier for viewers to locate using text-based search engines. Some video-hosting sites support hidden tags for website developers, allowing them to identify videos that may have been uploaded through their site to simplify later cataloging.

REVIEW AND CHECKLIST UPDATE

In this chapter, we first discussed the key attributes of Internet video and the basic system architecture. We next covered some of the unique applications of Internet video. This was followed by an overview of a typical system architecture and different platforms that can be used for Internet video services. We concluded with a discussion of some specific technologies, including client software and embedding videos that are hosted by other websites.

Chapter 15 Checklist Update

- ☐ Determine if video content is to be hosted on a private website or on a video-hosting website.
 - ☐ If on a private server, make sure there is adequate storage and bandwidth to deliver the video streams to all visitors. Also determine which delivery method and which compression method will be available to the desired audience.

- ☐ If on a video-hosting site, make sure that the video is properly formatted for the selected site and that the proper website coding has been done to be compatible with the host's APIs.
- ☐ If advertising is to be used, select the appropriate type:
 - ☐ Website or banner ads are suitable for almost any type of content.
 - ☐ Overlay ads for content that is to be closely associated with the content provider.
 - ☐ Preroll, mid-roll, or postroll ads for high-quality content where users will accept the interruptions.
- ☐ Consider using a CDN if large surges in audience are expected to occur or if the traffic levels do not warrant the costs of setting up a private server.
- ☐ If video content is to be inserted into a web page, consider using video websites such as YouTube to host the video, thereby reducing the costs of delivering the video files to website visitors.